

User's Guide



Notices

Carbonite Recover User's Guide, version 2.3.9, Thursday, May 18, 2023

If you need technical assistance, you can contact Customer Support. All basic configurations outlined in the online documentation will be supported through Customer Support. Assistance and support for advanced configurations may be referred to a Pre-Sales Systems Engineer or to Professional Services.

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, LLC.; 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, LLC. 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.

© 2023 Open Text. All rights reserved.

Contents

- Chapter 1 Carbonite Recover overview** **4**
 - How Carbonite Recover works 5
- Chapter 2 Requirements** **6**
 - Configuration and ports 12
- Chapter 3 Getting started** **13**
- Chapter 4 Carbonite Recover interface** **14**
- Chapter 5 Managing environments** **17**
 - Adding an environment 19
 - Viewing or editing environment details 20
- Chapter 6 Managing servers** **22**
 - Adding servers or existing target appliances 25
 - Creating a target appliance 28
- Chapter 7 Managing jobs** **31**
 - Creating a job 36
 - Testing failover 51
 - Failing over 61
 - Restoring 74
 - Failing back 85
 - Viewing job details 86
- Chapter 8 Configuring email notification** **90**
- Chapter 9 Viewing reports and company usage** **91**
- Chapter 10 Administration** **94**
 - Managing users 95
 - Adding a user 97
 - Viewing and editing user details 99
 - Managing workers 101
 - Installing a worker 103

Chapter 1 Carbonite Recover overview

Carbonite Recover protects any physical, virtual, or cloud server to the cloud. You identify the server you want to protect, and Carbonite Recover will replicate it to a virtual server stored in the cloud. The data is protected using Carbonite Availability real-time replication, also known as the Recover replication agent, which sends only file changes rather than copying an entire file, allowing you to more efficiently use server and network resources. In the event of a failure, you can failover to your replica virtual machine in the cloud with minimal downtime. See *How Carbonite Recover works* on page 5 for a workflow of the Carbonite Recover process.

How Carbonite Recover works

Begin with servers you want to protect in the cloud. These servers are referred to as source servers. Linux, and at least one worker. Carbonite will create a worker for you, but you must create the target appliance.

- **Target appliance**—A target appliance is a virtual server in the cloud that was created from a template provided by Carbonite. You must have at least one target appliance for Windows and one for Linux, and they can protect multiple source servers. However, you may need additional target appliances if you are protecting a larger number of disks or to help balance the load when protecting many servers. The target appliance maintains a replica of the data from the source servers you are protecting, and in the event of a failure, the data on the target appliance is used to quickly failover to a replica virtual machine in the cloud.

When protection begins, the target appliance is maintaining a replica of the data from the source servers you are protecting by using virtual hard disks attached to the target appliance.

In the event a source server fails, the worker quickly creates a replica virtual machine in the cloud and detaches the hard disks from the target appliance and attaches them to the new replica virtual machine.

You can run on the replica virtual machine in the cloud as long as needed. When you are ready, you can restore and failback from the replica virtual machine in the cloud back to your original server or to a different server, as needed.

Chapter 2 Requirements

Your environment must meet the following requirements.

- **Operating system**—The source servers you are protecting must be 64-bit architecture and must be one of the following Windows or Linux operating systems, with supported file system and kernel type.

Operating System	Version	File System	Kernel Type
Windows	2008 R2 Service Pack 1 or later	NTFS	Not applicable
	2012		
	2012 R2		
	2016	NTFS	
	2019		
Red Hat Enterprise Linux CentOS	6.8 through 6.10	Ext3	Default
	7.7 through 7.9	Ext4	
	8.1 through 8.3	XFS	
SUSE Linux Enterprise Server	12.2 through 12.4	Ext3 Ext4 XFS	Default
	15.0 through 15.2		
Ubuntu	16.04.2 through 16.04.4	Ext2 Ext3	Generic
	18.04.1 through 18.04.3	Ext4 XFS	



The following notes apply to Windows operating systems.

- A Windows server cannot be a Hyper-V server. Protection of a Hyper-V server is not supported.
- If your source is 2008 R2 Service Pack 1 or later, you must pre-install Microsoft .NET Framework version 4.5.1 or later before protecting the server.

The following notes apply to Linux operating systems.

- For all operating systems except Ubuntu, the kernel version must match the expected kernel for the specified release version. For example, if `/etc/redhat-release` declares the system to be a Redhat 7.7 system, the kernel that is installed must match that.

- Ubuntu networking must use traditional or NetworkManager. Ubuntu Netplan is not supported. You must disable Netplan and delete the /etc/netplan directory if you want to protect your Ubuntu server.
 - Stacking filesystems, like eCryptFS, are not supported.
-

- **Linux packages and services**—Each Linux source server must have the following packages and services installed before you can install and use Carbonite Recover. See your operating system documentation for details on these packages and utilities.
 - sshd (or the package that installs sshd)
 - lsb
 - parted
 - dmidecode
 - scp
 - which
 - libnsl (only required for Red Hat Enterprise Linux and CentOS version 8.0 and later)
- **SELinux policy**—SELinux should be disabled on your Linux source servers.
- **UEFI, trusted boot, secure boot**—The boot mode on your source servers cannot be UEFI (Unified Extensible Firmware Interface), trusted boot (tboot), secure boot, or other volume blocking mechanisms.
- **System memory**—At least 1 GB of memory is required on the source servers.
- **Server name**—The name of the source server must meet the following requirements.
 - Your source server name must be in ASCII format. Unicode file system support is included, it is the server name only that must be ASCII.
 - Due to VMware vCloud limitations, Windows server names can contain only letters, numbers, and hyphens. The name cannot be all numbers or start or end with a hyphen.
 - Due to VMware vCloud limitations, Linux server names can contain only letters, numbers, hyphens, and periods. The name cannot be all numbers or start or end with a hyphen or period.
 - All servers must have a unique server name.
- **VMware Tools**—All source servers hosted on VMware must have VMware Tools installed.
- **Windows Remote Management**—Carbonite Recover uses Windows Remote Management (WinRM) for several functions. For example, it uses it to push the necessary Recover replication agent software to your Windows source servers, to protect source disks that have more than 2 TB of data, and to determine the functional level of domain controllers. If you choose not to enable WinRM, functionality will be limited or you may have to take additional manual steps.

If you need to enable WinRM, type **winrm quickconfig** at a local command prompt on each server you are protecting and then type **y** to grant administrative rights remotely to local users.

- **Windows Remote Desktop**—If you have Network Level Authentication enabled on your source server, you will not be able to use Remote Desktop to access the replica virtual machine in the cloud after failover.
- **Windows caveats**—In some cases, there may be additional configurations that you might need for Windows sources. These configurations are difficult to identify because they depend

on unique situations, like whether you log in with local credentials or domain credentials. If you find issues with Carbonite Recover pushing the Recover replication agent software to your Windows source servers or inventorying the servers (gathering information from the servers), you may need one or more of the following changes on your Windows source servers.

- **Windows trusted hosts**—Add the worker to the trusted hosts list on the source so that WinRM communication is not blocked. (Administrative Templates, Windows Components, Windows Remote Management (WinRM), WinRM Client, Trusted Hosts)
- **Windows User Account Control (UAC)**—Disable User Account Control remote restrictions . Set HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System\LocalAccountTokenFilterPolicy to 1. You must reboot the server for the new setting to take effect.
- **Windows network type**—Update source server networking to be private. To change the network type, go to HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\NetworkList\Profiles. Look through the GUIDs and identify the NIC by the ProfileName. Then change Category to 1. You must reboot the server for the new setting to take effect.
- **SAN policy**—Some versions and editions of Windows set the default SAN policy to disabled. This keeps the non-boot volumes from mounting on the replica virtual machine after failover. To avoid this issue, enable your SAN policy on your source servers by using the following instructions.
 1. Open a command prompt on the source server.
 2. Type the following command.

```
diskpart
```
 3. Type the following command.

```
san
```
 4. If your policy is set to offline shared, change it using the following command.

```
san policy=onlineall
```
- **Clusters**—Clusters are not supported. Your source server cannot be in a cluster.
- **Domain servers**—If you are protecting domain joined servers, keep in mind the following.
 - You must include one or more domain controllers in your protection group so that the domain servers can resolve DNS and authenticate.
 - The primary domain controller must be authoritative.
 - The primary, authoritative domain controller should be a DNS server.
 - You must configure failover and failback server ordering.
 - The primary, authoritative domain controller must be the first server in the server ordering.
 - All remaining domain controllers must be after the primary, authoritative domain controller and before other domain servers.
 - Make sure you add the replica networking to Windows Sites and Services.
 - If you want to remotely install the replication agent on the servers, you must use domain credentials when adding the server to Carbonite Recover. If you do not use domain credentials, you must install the replication agent on the servers manually.

- **DNS**—You can configure your protection jobs to update DNS so that the source resolves to the replica virtual machine in the cloud after failover, and then back to the source (or an alternate server) after failback. However, this functionality is currently only available for Windows servers in a domain. Having only a DNS record is insufficient. The Windows source must have a computer account in the domain.

You should be protecting a source DNS server in order to be able to access other protected source servers after failover. Carbonite will provide you with an IP address to specify for your replica DNS server in the cloud. If you do not know the IP address to use, check with Carbonite.

Also, make sure you have added the target networking to Windows Sites and Services.

- **Automatic discovery and hypervisors**—Carbonite Recover can automatically discover source servers hosted on one of the following hypervisors. Servers not hosted on one of these hypervisors must be manually entered in Carbonite Recover.
 - **VMware**—You can use ESXi version 6.5 or later
 - **Hyper-V**—You can use any of the following Hyper-V versions.
 - Windows 2016
 - Windows 2012
 - Windows 2012 R2
 - Server Core 2008 R2 Service Pack 1 or later
 - Server Core 2012
 - Server Core 2012 R2
 - Hyper-V Server 2008 R2



Windows Management Instrumentation (WMI) must be enabled on your Hyper-V hosts for Carbonite Recover to discover hosted source servers.

- **Worker**—A worker server communicates with the Carbonite Recover infrastructure to receive and execute tasks on the source, target appliance, or in the cloud. Carbonite may create a worker for you, or if you are assigned the Administrator role for your user account, you can create a worker. If you create a worker, it must meet the following requirements.
 - **Operating system**—Windows 2016 or Windows 2019 (Standard or Datacenter)
 - **Memory**—At least 1 GB
 - **Processors**—At least 1 CPU
 - **Disk space**—At least 30 GB
 - **Windows .NET Framework**—Version 4.6.2 or later
- **Target appliance**—Carbonite Recover will walk you through creating a target appliance. You must configure the networking during the target appliance creation process so that it can communicate with the source servers you are protecting.

Keep in mind that a single virtual recovery appliance can protect a maximum of 10 sources or jobs with a maximum of 59 disks.

- **Networking**—Keep in mind the following networking caveats.





- You must establish a VPN between the source servers and the target appliance.
- If your source and target appliance are on different subnets you will need name resolution.
- When configuring your replica virtual machines, if you select a cloud network that is an isolated network, the servers that are failed over can communicate with each other but may not be able to communicate outside the isolated network. You have a few choices for working with this type of environment.
 - Include all required servers in the job group so all servers failover together.
 - Exclude a server from the job group, but add the required services from that server to another server that is in the job group.
 - Use an SSL VPN client to provide a route to the isolated network. This would be a point-to-site VPN so only the servers running the SSL VPN client will have connectivity.
- **Ports**—In order for your source servers and target appliance to communicate and transmit data, you must have specific ports open.
 - **Windows**—You must open ports 6320 and 6325 on Windows source servers and target appliances. You must also have Windows Remote Management (HTTP-In) configured. Port 5985 (for HTTP) and 5986 (for HTTPS) must be open for public and domain profiles. For source servers that are not on the same local network as the worker, confirm in Windows Remote Management (HTTP-In) that the **Scope** for **Remote IP address** is set to **Any IP address**.
 - **Linux**— You must open ports 1500, 6325, and 6326 on Linux source servers and target appliances. In order for Carbonite Recover to push the necessary Recover replication agent software to your Linux source servers, you must have port 22 open.
- **Time**—The clock on your source servers and your target appliance must be within a few minutes of each other, relative to UTC. Large time skews (more than five minutes) will cause Carbonite Recover errors.
- **Recovery points**—Carbonite Recover uses Microsoft and LVM technology for recovery point (snapshot) support. Keep in mind the following when using recovery points.
 - **Disk space**—If you use the default replica disk size when creating a job, a moderate amount of extra disk capacity is automatically added to accommodate recovery points. If you customize the disk size, no additional disk capacity is added. If you expect a high rate of data change, customize the disk size to increase the capacity for larger recovery points.
 - **Windows**—You can use VSS on your source servers for other uses outside Carbonite Recover, for example Microsoft Backup. Keep in mind that the driver for VSS is started before the Carbonite Recover replication agent driver. Therefore, if you use recovery points on your source servers and you revert any files on the source server, the Carbonite Recover replication agent will not be aware of the revert and the file change will not be replicated to the target appliance. The file change will be sent to the target appliance during the next synchronization process.
 - **Linux**—You can take recovery points of system volumes that are non-LVM as long as the root volume is logical, and you can take recovery points of data volumes managed under LVM. Recovery points of LVM data volumes will be created and stored on the target appliance also using LVM.

- **Web browser**—You require a web browser to access the Carbonite Recover web interface. A recent version of Google Chrome or Mozilla Firefox are the preferred browsers. You can also use other browsers such as Microsoft Internet Explorer version 11, however you may experience layout or appearance issues, such as field label misalignment. These issues should be display issues only and will not impact the functionality of your protection.

Configuration and ports

Your Carbonite Recover solution will consist of your source servers, at least one target appliance, and at least one worker. Carbonite will create the worker for you. See *Requirements* on page 6 for details on these components. You also require a web browser to access the Carbonite Recover interface.

Although data transferred between the source server and target appliance is encrypted using AES-256, a VPN is required between these servers.

Component to Component	Communication and Port	Arrow Color
Web Browser to Carbonite Recover	HTTPS port 443	
Source Server to Worker Target Appliance to Worker	HTTPS port 6326 HTTP port 5985 and HTTPS 5986	
Source Server to Target Appliance	Windows—Recover replication agent ports 6320, 6325, and 6326 Linux—Recover replication agent ports 1500, 6325, and 6326	
Worker to Source Server	Windows—SMB port 445 and Recover replication agent port 6325 Linux—HTTPS port 443	

Chapter 3 Getting started

Before you get started, make sure you have reviewed the Carbonite Recover *Requirements* on page 6. Then complete the following tasks, in order.

1. **Accept invitation**—Carbonite sends you an email invitation to access Carbonite Recover. See *Carbonite Recover interface* on page 14 for more details on accepting your invitation and an overview of the Carbonite Recover interface.
2. **Add environments** —You must create two environments, although one may already be created for you. An environment is a collection of servers. An environment may also have workers or a hypervisor host. A source environment is used to discover and protect your source servers. A target environment is used for the cloud and to provision resources and failover servers in the cloud. You must add a source environment that you will then populate with the source servers you want to protect. You must also add a target environment that you will then populate with your target appliance. See *Adding an environment* on page 19.



Carbonite may have created the target environment for you.

3. **Add servers to your source environment**—You must add servers to your source environment, either manually or through discovery. Discovery is the process of scanning a host to identify the servers on that host. See *Adding servers or existing target appliances* on page 25.
4. **Create a target appliance in your cloud environment**—Once you have created your target environment, you must create a target appliance in it. A target appliance is a virtual server in the cloud that was created from a template provided by Carbonite. You must have at least one target appliance for Windows and one for Linux, and they can protect multiple source servers. However, you may need additional target appliances if you are protecting a larger number of disks or to help balance the load when protecting many servers. The target appliance maintains a replica of the data from the source servers you are protecting, and in the event of a failure, the data on the target appliance is used to quickly failover to a replica virtual machine in the cloud. To create a target appliance, see *Creating a target appliance* on page 28. If you already have an existing target appliance, see *Adding servers or existing target appliances* on page 25.
5. **Create a job** —Once your source and target environments are prepared with sources and your target appliance, you can protect those sources. See *Creating a job* on page 36 for complete details.

Chapter 4 Carbonite Recover interface

- **First time access**—Carbonite will send you an email invitation to access Carbonite Recover. Click **Verify Email** to open the Carbonite Recover interface (hosted by Carbonite). Enter your first and last name. Enter and confirm a password. By continuing, you agree to the terms of service. Click **Confirm** to finalize your registration.
- **Bookmark**—Make sure that you bookmark the page so that you can easily return to the URL.
- **Companies**— Your Carbonite Recover account is associated with a Carbonite Recover company. Companies are created by Carbonite to determine which areas of the backend infrastructure users can access. You belong to a company based on company assignments. You can access the cloud resources assigned to your company, as well as any child companies that may exist under your assigned company. You can access your company or any child companies from the **View as company** drop-down list to the left of the bell notification icon. Select a company from the list or filter the list by typing in text and then selecting a company from the filter.
- **Dashboard**—The dashboard page displays each time you log in to Carbonite Recover. The dashboard provides a high-level overview and summarizes the status for the selected company and any of its child companies.

MACHINE	DISK SPACE USED	RECOVERY POINTS	LAST RECOVERY POINT	COMPANY
Beta	21.87% of 135.49 GB	1	04/17/2020 09:02:47	CompanyName
Alpha	18.95% of 135.49 GB	2	04/17/2020 15:02:02	CompanyName
Delta	12.07% of 135.49 GB	1	04/17/2020 09:02:47	CompanyName

- **Job Status**—This section updates dynamically as job status changes. You can click on the tiles in this section to view another table that separates the number of jobs in that state tile by the companies. You can click a company name hyperlink in that table to go to the **Jobs** page. (You will be asked to confirm if you are changing the currently viewed company.) A filter will automatically be applied to the **Jobs** page showing only the jobs that match the state tile you had selected. See *Managing jobs* on page 31 for details.

- **Disk Usage Summary**—The table shows the amount of disk space your source data is using on the replica virtual machine (attached to the target appliance until failover occurs). Note that if you used the default replica disk size when creating a job and enabled recovery points, a moderate amount of extra disk capacity was automatically added to the replica disk size to accommodate the recovery points. If you customized the disk size, no additional disk capacity was added.

Disk usage is collected once an hour, however this table is updated every five minutes to gather other data in the table. You will not see a source in the table until it has disk usage for at least one hour. Disk usage will not be updated during failover, restore, and failback. A server will be removed from the table if the job is deleted. Hover over the disk space to see the volumes and collection time for a specific source.

In the overflow menu on the right of a table row, you can select **View job details** to go the details page for that job. See *Viewing job details* on page 86 for details.

At the bottom of the table you will see the row numbers you are currently viewing and the total number of rows in the table. The paging buttons allow you to move between pages of the table. The single arrow buttons move forward or backward one page. The double arrow buttons move to the first or last page.

- **Console Manager**—Click any quick link in this section to jump to the corresponding page in the console.
- **Job Activity**—This section is updated dynamically as job activity occurs. The activity reported does not include manual interactions you may take, such as taking a recovery point. Click **View all activity** to jump to the **Jobs** page. See *Managing jobs* on page 31 for details.
- **Notifications**—You can access a list of notifications by clicking the bell icon near the upper right corner of any page. The color of the bell icon is customized depending on the type of notification. A green bell indicates that you have only information notifications. If you have one or more warning notifications, the bell will be yellow. If you have one or more error notifications, the bell will be red. The yellow warning color overrides green information, and red error overrides yellow warning. You can click the bell to view the notifications and the approximate times they were generated. You can dismiss individual notifications or all of the notifications in the current list. You can also filter the notification list to one type of notification by clicking on the circle for that notification type at the bottom of the notification list.
- **Gear icon**—If you click the gear icon in the upper right corner, you will find another menu of options.
 - **Manage instance**—This options allows you to view workers.
 - **Your name**—This option allows to control preferences specific to your account.



Users assigned the administrator role will be identified with (Admin) after their user name.

- **Preferences**—This tab allows you to set email notifications.
 - **Subject Prefix**—By default, the subject line of email alerts sent to your account email address will be prefaced with Carbonite Recover Notification

and the company name. This prefix allows you to recognize and filter emails specific to Carbonite Recover and companies. You can change or remove the first part of prefix as desired (not the company name). The remainder of the subject line will contain the notification content, truncated if necessary. The email body will contain the full notification content.

- **Notifications**—Select the type and level of notifications that you want to receive as email messages. If you do not select any type or level, you will not receive notifications as email messages. You will still get notifications in the Carbonite Recover web interface whether email notifications are enabled or disabled.
- **Security**—This table allows you to control the security of your account.
 - **Update Credentials**—You can enter your current password and then enter and confirm a new password. If you have forgotten your password and need to reset it, use the **Forgot password** link on the log in page.
 - **Two-Step Verification**—You can enable multi-factor authentication to provide more secure access to your Carbonite Recover account
 - **Get Started**—Click this link to enable multi-factor authentication. You will be asked to provide a phone number and to indicate if you want to receive a text or voice (text to speech) message. Enter the confirmation code that you receive to finalize multi-factor authentication.
 - **Change**—Click this link to change the phone number used for multi-factor authentication.
 - **Remove**—Click this link to remove multi-factor authentication.
- **Sign out**—This option immediately logs you out of the Carbonite Recover interface. If there is a period of inactivity, you will automatically be logged out. Any jobs that you have started will continue to run, even when you are logged out.
- **Support**—This option will open a new browser window to the Carbonite Support Knowledge Base
- **User's Guide**—This option will open a new browser window to the Carbonite Recover User's Guide.

Chapter 5 Managing environments

An environment is a collection of servers. An environment may also have workers or a hypervisor host. A source environment is used to discover and protect your source servers. A target environment is used for the cloud and to provision resources and failover servers in the cloud. You must add a source environment that you will then populate with the source servers you want to protect. You must also add a target environment that you will then populate with your target appliance.

When you create an environment, you will assign it a name and select the environment type. The following environment types are supported.

- **Google Cloud Platform**—This environment contains a Google Cloud Platform host, the source servers that you want to protect, and possibly a worker.
- **Microsoft Hyper-V**—This environment contains a host, the source that servers you want to protect, and possibly a worker.
- **VMware vSphere**—This environment contains a host, the source that servers you want to protect, and possibly a worker.
- **Custom**—This environment contains source that servers you want to protect. The servers could be in a hosted environment, but custom allows you to add the server without using the host. For example, you would have to use a custom environment for physical server or a server hosted in Microsoft Azure.
- **VMware vCloud**—This environment contains a cloud host, the target appliance, and a worker. This is your target cloud environment.

The **Environments** page provides high-level information and controls for your environments for the currently selected company.

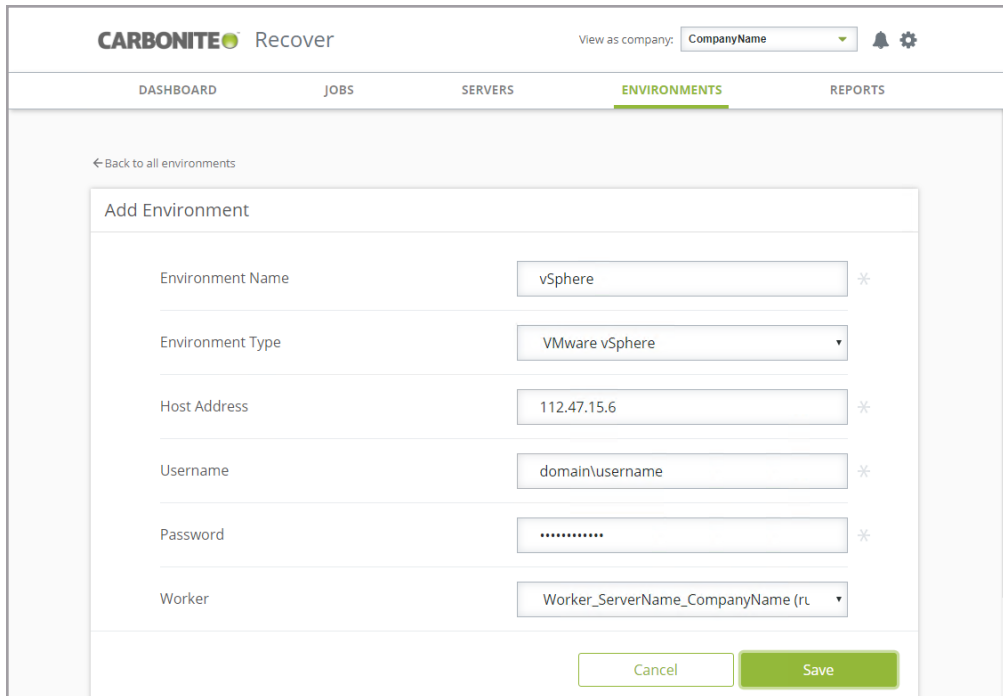
The following controls and status are available on the **Environments** page.

- **Add environment**—Click this button to add a new environment. See *Adding an environment* on page 19 for more details on this process.
- **Delete**—When at least one row in the table is selected, you can select this option in the overflow menu to the right of the add button to delete the selected environments. You cannot delete an environment that has servers in it.
- **Select All** and **Clear All**—Click the checkbox in the column heading to toggle between selecting all items on that page of the table or clearing all selections on that page of the table. This option will not apply to items on a page that are hidden by a search filter.
- **Sort**—Sort the table by clicking a column heading. When the arrow is pointing up, the table is sorted by that column in ascending order. When the arrow is pointing down, the table is sorted by that column in descending order.
- **Filter**—Text entered in a filter box and selected from a filter drop-down list will narrow the list displayed to only those rows that contain the search text and selected item.
- **Status**—The status indicates, by color and description, the health of the environment.
 - **Green**—A green circle indicates a good status.
 - **Yellow**—A yellow circle indicates a pending or warning status. Generally, Carbonite Recover is working, waiting on a pending process, or attempting to resolve the warning state.

- **Red**—A red circle indicates an error status. You will need to investigate and resolve the error.
- **Black**—A black circle indicates the status is unknown.
- **Table row overflow menu**—In the overflow menu on the right of a table row, you can select the following actions.
 - **View details**—Select this option to view or edit the environment details. See *Viewing or editing environment details* on page 20 for more details.
 - **Add server**—Select this option to add servers to this environment. For hypervisor environments, Carbonite Recover will discover the servers on the host, and you can select which servers you want to add. For custom environments, you will need to specify the servers by IP address. See *Adding servers or existing target appliances* on page 25 for details.
 - **Delete**—Select this option to delete the environment. You cannot delete an environment that has servers in it.
- **Table paging**—At the bottom of the table you will see the row numbers you are currently viewing and the total number of table rows. Paging buttons allow you to move between pages of the table. The single arrow buttons move forward or backward one page. The double arrow buttons move to the first or last page.

Adding an environment

1. On the **Environments** tab, click **Add environment**.
2. Identify your environment.

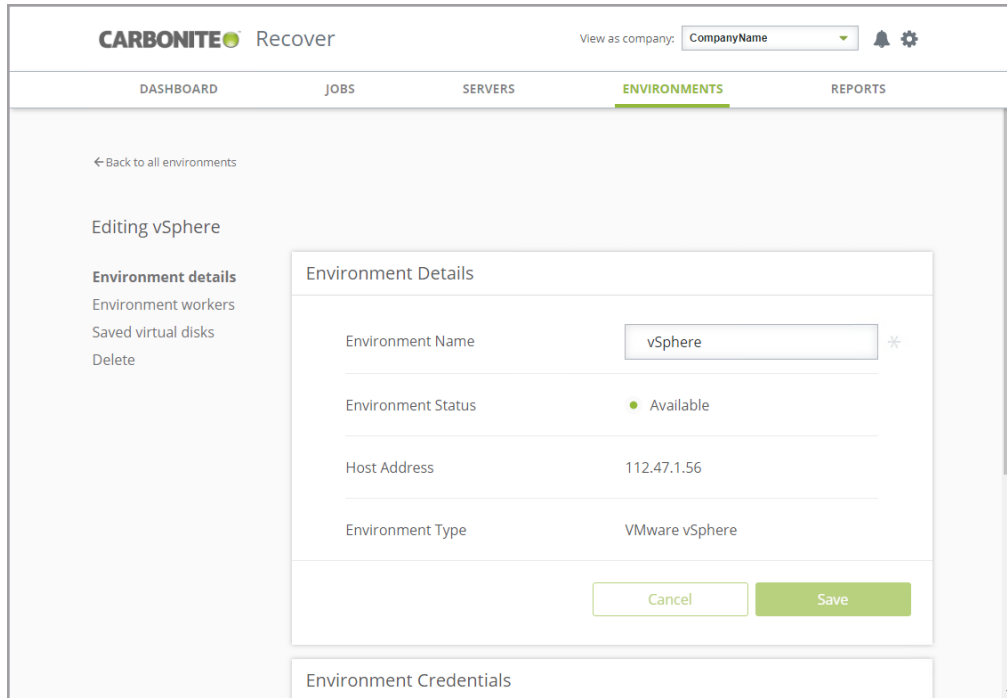


The screenshot shows the Carbonite Recover web interface. At the top, there is a navigation bar with the Carbonite Recover logo, a 'View as company:' dropdown menu set to 'CompanyName', and notification and settings icons. Below this is a secondary navigation bar with tabs for 'DASHBOARD', 'JOBS', 'SERVERS', 'ENVIRONMENTS' (which is highlighted), and 'REPORTS'. The main content area features a 'Back to all environments' link and a 'Add Environment' form. The form contains the following fields: 'Environment Name' (text input with 'vSphere' and an asterisk), 'Environment Type' (dropdown menu with 'VMware vSphere'), 'Host Address' (text input with '112.47.15.6' and an asterisk), 'Username' (text input with 'domain\username' and an asterisk), 'Password' (password input with asterisks and an asterisk), and 'Worker' (dropdown menu with 'Worker_ServerName_CompanyName (r...' and an asterisk). At the bottom of the form are 'Cancel' and 'Save' buttons.

- Environment Name—Specify a unique name for this environment that will distinguish it from other Carbonite Recover environments.
 - Environment Type—Select the type of environment you want to add.
 - Host Address—For Microsoft Hyper-V and VMware vSphere environments, specify the IP address for the host.
 - Host Address—For Google Cloud Platform and Microsoft Hyper-V environments, specify the IP address for the host.
 - Project name—For Google Cloud Platform, type the name of the project associated with the Google Cloud Platform environment.
 - Username—For Microsoft Hyper-V and VMware vSphere environments, specify a user name with access to the host. Your vCloud target environment should already be created for you. If it is not, create it and specify the user name provided by Carbonite.
 - Username—For Google Cloud Platform and Microsoft Hyper-V, specify a user name with access to the host. Your vCloud target environment should already be created for you. If it is not, create it and specify the user name provided by Carbonite.
 - Password—Specify the password associated with the user you have entered.
 - Organization—For VMware vCloud environments specify the vCloud organization where your servers will be protected. If you do not know which organization to select, contact Carbonite. If you have only been granted access to one organization or you have specified a different environment type, you will not see this option.
 - Worker—Select a worker for this environment.
3. When you have identified your environment, click **Save**.

Viewing or editing environment details

1. On the **Environments** page, find the table row of the environment you want to view or edit. In the overflow menu for that table row, click **View details**.
2. On the **Environment details** tab, view or modify the environment details in the top section.



- **Environment name**—Change the name of the environment, if required.
 - **Environment Status**—The status indicates, by color and description, the health of the environment.
 - **Green**—A green circle indicates a good status.
 - **Yellow**—A yellow circle indicates a pending or warning status. Generally, Carbonite Recover is working, waiting on a pending process, or attempting to resolve the warning state.
 - **Red**—A red circle indicates an error status. You will need to investigate and resolve the error.
 - **Black**—A black circle indicates the status is unknown.
 - **Host address**—If the environment has a hypervisor host, this is a read-only field that displays the address of the host.
 - **Organization**—If the environment is vCloud, this is a read-only field that displays the vCloud organization.
 - **Environment type**—This is a read-only field that displays the type of environment.
 - **Project ID**—If the environment is Google Cloud Platform, this is a read-only field that displays the name of the project associated with the Google Cloud Platform environment.
3. If you have changed the name of the environment, click **Save**.

4. On the same **Environment details** tab, if the environment has a hypervisor host, you can modify the credentials used to access the host. Custom environments do not have credentials. Be sure to click **Save** to save any changes to the credentials.
5. On the **Environment workers** tab, you can see the worker the environment is using and its status. You cannot make any changes on this tab. If you need to change workers, you will have to delete the environment and re-create it. You cannot delete an environment that has servers in it.
6. If you are viewing the details of your target vCloud environment, you will see a **Saved virtual disks** tab. This tab shows the servers and their disks that you have saved from previous protections. Expand the server row to see the saved disks for that server.
7. On the **Delete** tab, you can click **Delete** to delete the environment. You cannot delete an environment that has servers in it.

Chapter 6 Managing servers

Once you have an environment created, you can add servers or target appliances to the environment manually or through discovery. Discovery is the process of scanning a host in an environment to identify the servers on that host. You can also create a target appliance in your target cloud environment.

The **Servers** page provides high-level information and controls for your servers for the currently selected company.

<input type="checkbox"/>	NAME ^	ENVIRONMENT	STATUS	NETWORK ID	AGENT VERSION	OPERATING SYSTEM	
<input type="checkbox"/>	Alpha	vSphere	Unprotected	10.6.7.41	8.4.0.198.0	Windows Server 2012 ...	⋮
<input type="checkbox"/>	Beta	vSphere	Unprotected	10.6.7.66	8.4.0.198.0	Windows Server 2012 ...	⋮
<input type="checkbox"/>	Windows appliance	vCloud	Inactive target appliance	10.6.7.67	8.4.0.198.0	Windows Server 2012 ...	⋮

The following controls and status are available on the **Servers** page.

- **Add server**—Click this button to add servers to an environment. See *Adding servers or existing target appliances* on page 25 for more details on adding servers or target appliances that have already been created. If you need to create a target appliance, see *Creating a target appliance* on page 28.
- **Table overflow menu**—When at least one row in the table is selected, you will have an overflow menu to the right of the add button.
 - **Refresh**—Select this option to re-inventory the selected servers and gather information from them.
 - **Update credentials**—Select this option to update the credentials used to access the selected servers.
- **Select All** and **Clear All**—Click the checkbox in the column heading to toggle between selecting all items on that page of the table or clearing all selections on that page of the table. This option will not apply to items on a page that are hidden by a search filter.

- **Sort**—Sort the table by clicking a column heading. When the arrow is pointing up, the table is sorted by that column in ascending order. When the arrow is pointing down, the table is sorted by that column in descending order.
- **Filter**—Text entered in a filter box and selected from a filter drop-down list will narrow the list displayed to only those rows that contain the search text and selected item.
- **Status**—The status indicates, by color and description, the health of the server. If you have a job associated with the server, the status will show the job health, unless there is a problem with the server. In that case, the status of the server will be displayed. If there are no jobs associated with the server, the status is always the server health.
 - **Green**—A green circle indicates a good status.
 - **Yellow**—A yellow circle indicates a pending or warning status. Generally, Carbonite Recover is working, waiting on a pending process, or attempting to resolve the warning state.
 - **Red**—A red circle indicates an error status. You will need to investigate and resolve the error.
 - **Gray**—A gray circle indicates an unprotected server.
 - **Black**—A black circle indicates the status is unknown.



If your server is identified as invalid or unsupported, review the *Requirements* on page 6.

- **Table row overflow menu**—In the overflow menu on the right of a table row, you can select the following actions. The available actions will change depending on the role and state of the server.
 - **Refresh**—Select this option to re-inventory the server and gather information from it.
 - **Restart**—Select this option to restart a failed target appliance creation.
 - **Protect**—Select this option to protect a source server. This option is only available for source servers in an unprotected state. See *Creating a job* on page 36 for more details on this process.
 - **Install replication agent**—Select this option to install the Recover replication agent on the server. The replication agent is the engine that powers synchronization and replication from your source server to your target appliance in the cloud and from your replica virtual machine in the cloud to your failback target during restoration. This is an optional installation because the protection process will automatically install the replication agent during appliance creation and on the source during job creation, if it is not already installed. Keep in mind that port 22 must be open on Linux servers. On Windows servers, WinRM is required to push the installation. See *Requirements* on page 6 for details.
 - **Upgrade replication agent**—Select this option to automatically upgrade the replication agent on an unprotected server or a target appliance. You can also upgrade a protected source server if the target appliance is already upgraded. If the target appliance for the job has not been upgraded, then you cannot upgrade that protected source server. You can go to the **Jobs** page to upgrade all servers in a job together. If your target appliance

(active or inactive) is running a later version of the replication agent, you will not be able to protect an unprotected server running an older version of the replication agent. In this case, upgrade the unprotected server before creating a protection job. Keep in mind that port 22 must be open on Linux servers. On Windows servers, WinRM is required to push the installation. See *Requirements* on page 6 for details.

- **Reconcile license**—If your server is a target appliance and the replication agent license version does not match the replication agent software version, you will have this option to manually update the license version to match the software version. This option also appears for source servers that are part of a job. If your source server is not part of a job, the license version will automatically be updated when the server is used in a job.
- **Update credentials**—Select this option to update the credentials used to access the server.
- **Gather diagnostics**—Select this option to collect log files from the replication agent. Because this process is gathering several pieces of information across the network, it may take several minutes to complete the gathering process.
- **Download diagnostics**—Select this option to download the last collected diagnostics log file to the local machine.
- **Configuration**—This option is not currently supported. Do not select it or modify any settings on the associated screens or your protections will fail.
- **Remove**—Select this option to remove the server from the list. You cannot remove a server that has an established job.
- **Shut down (soft)**—Select this option to shut down the guest operating system gracefully. This option is only available for servers running in your target cloud environment.
- **Power off (hard)**—Select this option to abruptly power off the server without waiting for the guest operating system to shut down gracefully. This is like turning off the power switch. This option is only available for servers running in your target cloud environment.
- **Power on**—Select this option to power on the server. This option is only available for servers running in your target cloud environment.
- **Reset (hard)**—Select this option to reboot the server without waiting for the guest operating system to shut down gracefully. This option is only available for servers running in your target cloud environment.
- **Delete VM**—Select this option to delete the server from your target cloud environment. This option is only available for servers running in your target cloud environment.
- **Remote Console**—Select this option to access the server console remotely. The console session will open in a new browser tab or window depending on your browser configuration. This option is only available for servers running in your target cloud environment.
- **Table paging**—At the bottom of the table you will see the row numbers you are currently viewing and the total number of table rows. Paging buttons allow you to move between pages of the table. The single arrow buttons move forward or backward one page. The double arrow buttons move to the first or last page.

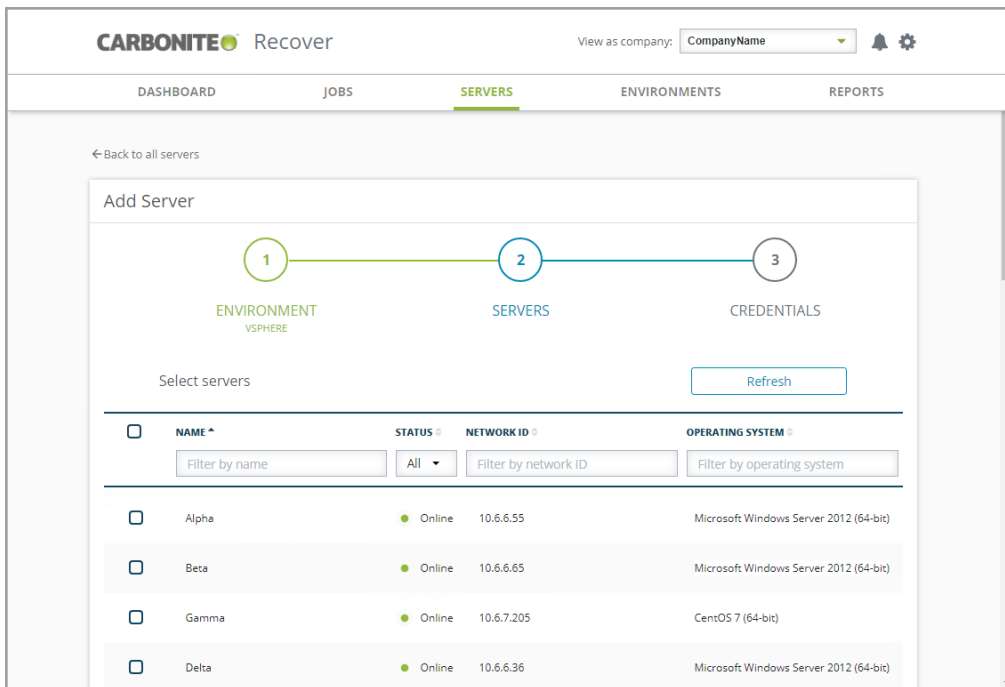
Adding servers or existing target appliances

When adding servers from a hypervisor environment, Carbonite Recover will discover the servers on the host, and you can select which servers you want to add. For custom environments, you will need to specify the servers by IP address. You can also add an existing target appliance to a vCloud environment. However, if you need to create a target appliance, see *Creating a target appliance* on page 28.

1. Create an environment, if you do not have one already. See *Adding an environment* on page 19.
2. On the **Servers** tab, click **Add server**. You can also select **Add server** from the overflow menu in a row on the **Environments** tab.
3. On the **Add Server** page, confirm or select your desired environment. If you have selected your cloud environment, make sure you have selected **Discover servers**.
4. Click **Next**. The next step will vary depending on if you have a hosted or custom environment.

Hosted environment

1. Select the servers or target appliances you want to add to your environment from the discovery of servers on your host by clicking a checkmark in the checkbox to the left of a server name.



The screenshot shows the Carbonite Recover interface. At the top, there's a navigation bar with 'DASHBOARD', 'JOBS', 'SERVERS' (highlighted), 'ENVIRONMENTS', and 'REPORTS'. Below this, a 'Back to all servers' link is visible. The main content area is titled 'Add Server' and features a progress indicator with three steps: 1. ENVIRONMENT (VSPHERE), 2. SERVERS, and 3. CREDENTIALS. Under the 'SERVERS' step, there's a 'Select servers' section with a 'Refresh' button. Below this is a table with columns for NAME, STATUS, NETWORK ID, and OPERATING SYSTEM. The table contains four rows of server information.

<input type="checkbox"/>	NAME	STATUS	NETWORK ID	OPERATING SYSTEM
<input type="checkbox"/>	Alpha	Online	10.6.6.55	Microsoft Windows Server 2012 (64-bit)
<input type="checkbox"/>	Beta	Online	10.6.6.65	Microsoft Windows Server 2012 (64-bit)
<input type="checkbox"/>	Gamma	Online	10.6.7.205	CentOS 7 (64-bit)
<input type="checkbox"/>	Delta	Online	10.6.6.36	Microsoft Windows Server 2012 (64-bit)



You have the following table controls available on step 2 of the **Add Server** workflow.

- **Refresh**—Click this button to rescan the host and refresh the list of servers.
- **Table checkbox column heading**—Use the checkbox column heading to select or deselect only the rows that are visible on the current page. If you want to select or deselect all rows on all pages, use the **Select All** or **Clear All Selections** links.
- **Sort**—Sort the table by clicking a column heading. When the arrow is pointing up, the table is sorted by that column in ascending order. When the arrow is pointing down, the table is sorted by that column in descending order.
- **Filter**—Text entered in a filter box and selected from a filter drop-down list will narrow the list displayed to only those rows that contain the search text and selected item.
- **Table paging**—At the bottom of the table you will see the row numbers you are currently viewing and the total number of table rows. Paging buttons allow you to move between pages of the table. The single arrow buttons move forward or backward one page. The double arrow buttons move to the first or last page.

You cannot protect an offline server.

Your servers must have at least one NIC attached to the server in order for the server to be discoverable.

Make sure all enabled NICs are operational. An enabled but unplugged NIC will cause a server to fail to be added to your environment. Disabled NICs are not an issue.

2. Click **Next** to continue.
 3. If your environment type is the target cloud, you need to select the server as an **Appliance**, meaning it is the target appliance that will protect your source servers. Only leave the checkbox disabled if the server is a replica virtual machine (a source that has already failed over to the cloud.) Click **Next** to continue.
 4. Specify credentials for Carbonite Recover to use to access the servers or target appliances that you are adding.
-



For Windows source servers that are part of a domain, you need to use domain credentials if you want to remotely install the replication agent on the servers. If you do not use domain credentials, you will need to install the replication agent on the servers manually.

For Linux source servers, if you choose to use a non-root user, it must be a user with sudo permissions because Carbonite Recover needs super user privileges.

5. Click **Finish** to add the servers or target appliances to the hosted environment.

Custom environment

1. Specify the network name or IP address of the server you want to add along with credentials for Carbonite Recover to use to access the servers.



For Windows source servers that are part of a domain, you need to use domain credentials if you want to remotely install the replication agent on the servers. If you do not use domain credentials, you will need to install the replication agent on the servers manually.

For Linux source servers, if you choose to use a non-root user, it must be a user with sudo permissions because Carbonite Recover needs super user privileges.

#	NETWORK ID	USERNAME	PASSWORD
1	172.31.206.200	administrator
2	172.31.206.201	administrator

2. Click **Add Server** to add another row to the table. If you need to remove a server from the list, click the overflow menu at the right side of a table row and select **Remove**.
3. Click **Finish** to add the servers to the custom environment.



Servers that show invalid credentials may actually be unreachable. Try to refresh the server or updating your credentials. If that does not work, confirm the server is reachable and then remove and re-add the server.

Creating a target appliance

A target appliance is a virtual server in the cloud that was created from a template provided by Carbonite. You must have at least one target appliance for Windows and one for Linux, and they can protect multiple source servers. However, you may need additional target appliances if you are protecting a larger number of disks or to help balance the load when protecting many servers.

The target appliance maintains a replica of the data from the source servers you are protecting, and in the event of a failure, the data on the target appliance is used to quickly failover to a replica virtual machine in the cloud.

Carbonite Recover will walk you through creating a target appliance. You must configure the networking during the target appliance creation process so that it can communicate with the source servers you are protecting.

Use these instructions to create a target appliance. (If you need to add an existing target appliance to your environment, see *Adding servers or existing target appliances* on page 25.)

1. Create a target cloud environment, if you do not have one already. See *Adding an environment* on page 19.
2. On the **Servers** tab, click **Add server**. You can also select **Add server** from the overflow menu in the target cloud row on the **Environments** tab.
3. On the **Add Server** page, confirm your target cloud environment is selected and then select **Create a new appliance**.
4. Click **Next** to continue.
5. Specify the storage and platform for the target appliance.

The screenshot shows the Carbonite Recover interface. At the top, there's a navigation bar with 'DASHBOARD', 'JOBS', 'SERVERS' (highlighted), 'ENVIRONMENTS', and 'REPORTS'. Below this is a 'Back to all servers' link. The main content area is titled 'Add Server' and features a progress indicator with three steps: 1. ENVIRONMENT (vCLOUD), 2. STORAGE, and 3. CONFIGURATION. The 'STORAGE' step is currently active. Below the progress bar, there are three dropdown menus: 'Storage Policy' set to 'vCloud Storage Policy', 'Organization vDC' set to 'Org_vDC', and 'Platform' set to 'Windows'. At the bottom, there are three buttons: 'Cancel', 'Back', and 'Next'.

- **Storage Policy**—Select the storage policy to use from the cloud environment you selected. If you do not know your storage policies, check with Carbonite.

- **Organization vDC**—Select the organization vDC to use from the storage policy selected. If you do not know your organization vDC, check with Carbonite.
 - **Platform**—Your target appliance operating system must be the same as your source servers. Therefore, if you are protecting Windows servers, select the **Windows** platform. If you are protecting Linux servers, select the **Linux** platform.
6. Click **Next** to continue.
 7. Create a target appliance by specifying the following options.

The screenshot shows the Carbonite Recover interface. At the top, there's a navigation bar with 'DASHBOARD', 'JOBS', 'SERVERS' (highlighted), 'ENVIRONMENTS', and 'REPORTS'. Below this is a 'Back to all servers' link. The main content area is titled 'Add Server' and features a three-step process diagram: 1. ENVIRONMENT (vCLOUD), 2. STORAGE (vCLOUD STORAGE POLICY), and 3. CONFIGURATION. The CONFIGURATION step is active, showing a form with the following fields:

- Name:** A text input field containing 'ApplianceName' with an asterisk indicating it's required. Below it is a checked checkbox labeled 'Use appliance name as host name'.
- Container:** A dropdown menu with 'vCloud_vApp' selected and an asterisk.
- Size:** A dropdown menu with 'Default Size' selected and an asterisk.
- Network:** A dropdown menu with 'Org_vDC_Network' selected and an asterisk.

- **Name**—Specify the virtual machine display name for the target appliance. Keep in mind the following when specifying a host name.
 - If you have enabled Use appliance name as host name, this field will be limited to 15 characters because the guest name cannot exceed 15 characters.
 - Due to VMware vCloud limitations, Windows server names can contain only letters, numbers, and hyphens. The name cannot be all numbers or start or end with a hyphen.
 - Due to VMware vCloud limitations, Linux server names can contain only letters, numbers, hyphens, and periods. The name cannot be all numbers or start or end with a hyphen or period.
 - Due to Google Cloud Platform limitations, both Windows and Linux server names must follow these guidelines:
 - The first character must be a lowercase letter.
 - All of the following characters (except for the last character) must be a dash, lowercase letter, or digit.
 - The last character must be a lowercase letter or digit.
- **Host Name**—If you have **Use appliance name as host name** disabled, you will see this field. Specify the guest name for the target appliance. If you want it to be the same as

the display name, enable **Use appliance name as host name**. Keep in mind the following when specifying a host name.

- The name cannot exceed 15 characters.
- Due to VMware vCloud limitations, Windows server names can contain only letters, numbers, and hyphens. The name cannot be all numbers or start or end with a hyphen.
- Due to VMware vCloud limitations, Linux server names can contain only letters, numbers, hyphens, and periods. The name cannot be all numbers or start or end with a hyphen or period.
- **Container**—Select or create a container where you want to create the target appliance.
- **Container Name**—If you are creating a new container, specify the name. If the name you enter already exists, Carbonite Recover will append a unique number to the name.
- **Size**—Select the size of the target appliance.
 - **Default**—This is option will create an appliance using the preferred Carbonite Recover default.
 - **Predetermined sizes**—You may or may not have other predefined sizes to select from. If you have them, you can select an alternate size. If you are uncertain what the specifications are for the size, contact Carbonite.
- **Network**—Select the network that you want the target appliance to use. If you do not know your available networks, check with Carbonite.
- **Adapter**—Select a network adapter. The types available in the list will depend on the operating system you have selected.
- **IP Mode**—Select **Pool** if you want the target appliance to be assigned an IP address from a pool of addresses. Select **Manual** and specify an **IP Address** if you want to assign a specific IP address to the adapter.

8. Click **Finish** to create the new target appliance.



It may take a minute or two for the target appliance to appear on the **Servers** page. Additionally, it will take time for the target appliance to finish creation, for example 15-30 minutes. Target appliance creation time is dependent on the Carbonite hardware and your connection to the cloud.

The target appliance will automatically power on during creation.

If the Recover replication agent version on the target appliance is not displayed on the **Servers** page after target appliance creation has completed, re-enter the target appliance credentials from the **Servers** page to re-inventory the target appliance, which will pull the Recover replication agent version number from the target appliance.

When creating a Linux target appliance, all file system packages supported by Carbonite Recover will be installed in order to properly format disks during protection and failover.

Chapter 7 Managing jobs

You protect source servers in groups, even if you have just one source in a group. Groups allow you to manage your servers in orchestration with each other.

The **Jobs** page provides high-level information and controls for your jobs for the currently selected company. You can view the individual servers in the group by clicking the right arrow to open the drop-down area below the group. Click the down arrow to close the drop-down area.



Within the expanded group, the servers displayed under the **Source** and **Target** headings changes depending on where you are in your job lifecycle.

- **Protecting and failover**—During the protecting and failover states, the source of the job is your source server and the target of the job is your target appliance.
- **Restoring and fallback**—During the restoring and fallback states, the source of the job is your replica virtual machine in the cloud and the target of the job is your fallback target.

Jobs for CompanyName

GROUP NAME ^ TARGET ENVIRONMENT ▾ JOBS ▾ STATUS ▾

Filter by group name All All

Alpha and Beta vCloud 2 Synchronizing

SOURCE	TARGET	STATUS
Alpha	WindowsAppliance	Synchronizing (60.9%)
Beta	WindowsAppliance	Synchronizing (88.3%)

Statistics

- Disk Queue -
- Initial Mirror Complete False
- SSH Direct Connection False
- Mirror Remaining 14.4 GB
- Mirror Skipped 5.4 GB
- Recovery Point Latency 0 seconds
- Replication Queue -
- Data Sent 36.4 GB
- Compressed Data Sent 36.4 GB

1 - 1 of 1

©2020. All rights reserved. | Support | Terms of Service | Privacy

The following controls are available on the **Jobs** page.

- **Table overflow menu**—When at least one row in the table is selected, you will have an overflow menu at the top of the **Jobs** page. Select **Delete** to delete the selected job groups. You will be prompted if you want to keep the disks used in the job and reuse them in later protections. In some cases, like after failover, reusing the disks is not available because the disks have already been used with the replica virtual machine. If you can reuse disks, it makes synchronization faster in the future because only different files need to be synchronized. If you

want to reuse disks, you must reuse all of the same disks and the server must exist within the same environment created in Carbonite Recover. Previous recovery points will not be available when reusing disks after deleting.

- **Select All and Clear All**—Click the checkbox in the column heading to toggle between selecting all items on that page of the table or clearing all selections on that page of the table. This option will not apply to items on a page that are hidden by a search filter.
- **Sort**—Sort the table by clicking a column heading. When the arrow is pointing up, the table is sorted by that column in ascending order. When the arrow is pointing down, the table is sorted by that column in descending order.
- **Filter**—Text entered in a filter box and selected from a filter drop-down list will narrow the list displayed to only those rows that contain the search text and selected item.
- **Expand and collapse**—Click the right arrow to expand a job group to see high-level details. Click the down arrow to collapse the details.
- **Hyperlink control**—Click the name of a job group to see additional details for the group and the individual jobs. See *Viewing job details* on page 86.
- **Status**—The status indicates, by color and description, the health of your jobs both at the group level and the individual job level.
 - **Green**—A green circle indicates a good status.
 - **Yellow**—A yellow circle indicates a pending or warning status. Generally, Carbonite Recover is working, waiting on a pending process, or attempting to resolve the warning state.
 - **Red**—A red circle indicates an error status. You will need to investigate and resolve the error.
 - **Black**—A black circle indicates the status is unknown.



If an arrow and a number display next to the status for an individual job, you can click the arrow to view additional status messages. These additional statuses are from the replication agent and can provide further information when a job is in an error state.

- **Table row overflow menu**—In the overflow menu on the right of a table row, you can select the following actions for job groups or the individual servers within a group. The available actions will change depending on the state of the job. Group actions will only be available when all servers in the group can safely perform that action. If you have only one server in a group, you will only have group actions.
 - **View details**—This option is only available at the group level. It is not available in the overflow menu for individual jobs within a group. Select this option to see details for the group and the individual jobs. See *Viewing job details* on page 86.
 - **Upgrade replication agents**—Select this option to automatically upgrade the replication agent on the servers in the group. Carbonite Recover will handle the required order of upgrading the target appliance first and then the source servers. This option is only available for jobs that are synchronizing to the target appliance in the cloud or protecting to the target appliance in the cloud. You will be blocked from upgrading jobs that have failed over or are being restored. This option is not available when the servers

are already running the latest replication agent version. Keep in mind that port 22 must be open on Linux servers. On Windows servers, WinRM is required to push the installation. See *Requirements* on page 6 for details.

- **Add recovery point**—Select this option to manually add a recovery point, and then confirm you want to add the recovery point. Keep in mind, scheduled recovery points will continue to be taken, if configured. Also note that recovery points will be skipped if the initial synchronization is not yet complete.
- **Failover**—Select this option to begin live or recovery point failover. See *Failing over* on page 61 for more details on this process.
- **Test failover**—Select this option to begin test failover. See *Testing failover* on page 51 for more details on this process.
- **Restore**—Select this option to begin restoration. This process is for jobs that have already been failed over to the cloud. It takes the replica virtual machine in the cloud and restores it back to your original source or another server. See *Restoring* on page 74 for more details on this process.
- **Failback**—Select this option after restoration is complete, to finalize the failback process. See *Failing back* on page 85 for more details on this process.
- **Reprotect**—Select this option to restart a job after you have restored and failed back. You will be prompted to reuse the existing disks from the last job or you can create new disks. In either case, the replica virtual machine created in the cloud from the last job will be deleted. Previous recovery points will be available when reusing disks after reprotecting.
- **Start**—Select this option to start a paused or stopped job. You can also use this action to restart the last failed operation. The restart action will restart all failed operations within the job group.
- **Stop**—Select this option to stop the job. Data changes will not queue on the source (if you are protecting) or replica virtual machine (if you are restoring). Data synchronization will restart from the beginning when the job is restarted.
- **Pause**—Select this option to pause the job. Data changes will queue on the source (if you are protecting) or replica virtual machine (if you are restoring). The changes will be transmitted once the job is resumed.
- **Undo test failover**—Select this option to undo a test failover. The replica virtual machine created in the cloud will be deleted and the job will be restarted.
- **Update DNS**— If you enabled DNS updates, you can trigger the DNS update process with this option. Edit your DNS settings if needed, and then click **Update DNS**. This option is only available for jobs with Windows source servers.
- **Configuration**—This option is only available at the group level. Select this option to view and, if desired, to edit the job configuration. A job can only be edited at certain stages of the job lifecycle. Also, not all settings can be edited. If want to edit any settings, click **Edit job group**, make the changes on any tab, and then click **Verify** to go through the validation process to confirm the new settings are acceptable. For details on the various job settings, see the descriptions used when creating the protection in *Creating a job* on page 36.
- **Delete**—Select this option to delete the job. This option may not be available depending on the state of your job. You will be prompted if you want to keep the disks used in the job and reuse them in later protections. In some cases, like after failover, reusing the disks is

not available because the disks have already been used with the replica virtual machine. If you can reuse disks, it makes synchronization faster in the future because only different files need to be synchronized. If you want to reuse disks, you must reuse all of the same disks and the server must exist within the same environment created in Carbonite Recover. Previous recovery points will not be available when reusing disks after deleting. Depending on the state of your job, you may also be prompted if you want to keep the replica virtual machine in the cloud. If you delete the disks, you will not be able to keep the replica virtual machine.

- **Statistics**—These statistics are cumulative for all of the jobs in the group.
 - **Disk Queue**—This is the amount of disk space being used to queue data on the source servers (when protecting) or on the replica virtual machines (when restoring).
 - **Initial Mirror Complete**—This field indicates if all of the initial copies of data have completed from your source servers to the target appliances (when protecting) or from the replica virtual machines to the failback targets (when restoring).
 - **SSH Direct Connection**—This field will always be false.
 - **Mirror Remaining**—This is the amount of data remaining to be sent from the source servers to the target appliances (when protecting) or from the replica virtual machines to the failback targets (when restoring).
 - **Mirror Skipped**—This is the amount of data that has been skipped because the data is not different on the source servers and target appliances (when protecting) or on the replica virtual machines and failback targets (when restoring).
 - **Recovery Point Latency**—This is the longest length of time replication is behind on any one target appliance compared to the source server they are protecting or on any one failback target compared to the replica virtual machine they are restoring from. This is the longest 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 synchronization data. If you are synchronizing and failover (or synchronizing and failback), the data on the target appliance (or the failback target) will be at least as far behind as the replication point latency. It could potentially be further behind depending on the circumstances of the synchronization. If synchronization is idle and you failover (or failback), the data will only be as far behind as the replication point latency time.
 - **Replication Queue**—This is the amount of disk space being used to queue replication data on the source servers (when protecting) or replica virtual machines (when restoring).
 - **Data Sent**—This is the total amount of data sent from the source servers to the target appliances (when protecting) or from the replica virtual machines to the failback targets (when restoring).
 - **Compressed Data Sent**—This is the total amount of compressed data sent from the sources servers to the target appliances (when protecting) or from the replica virtual machines to the failback targets (when restoring). If compression is disabled, this statistic will be the same as bytes sent.
- **Table paging**—At the bottom of the table you will see the row numbers you are currently viewing and the total number of table rows. Paging buttons allow you to move between pages of the table. The single arrow buttons move forward or backward one page. The double arrow

buttons move to the first or last page.

Creating a job

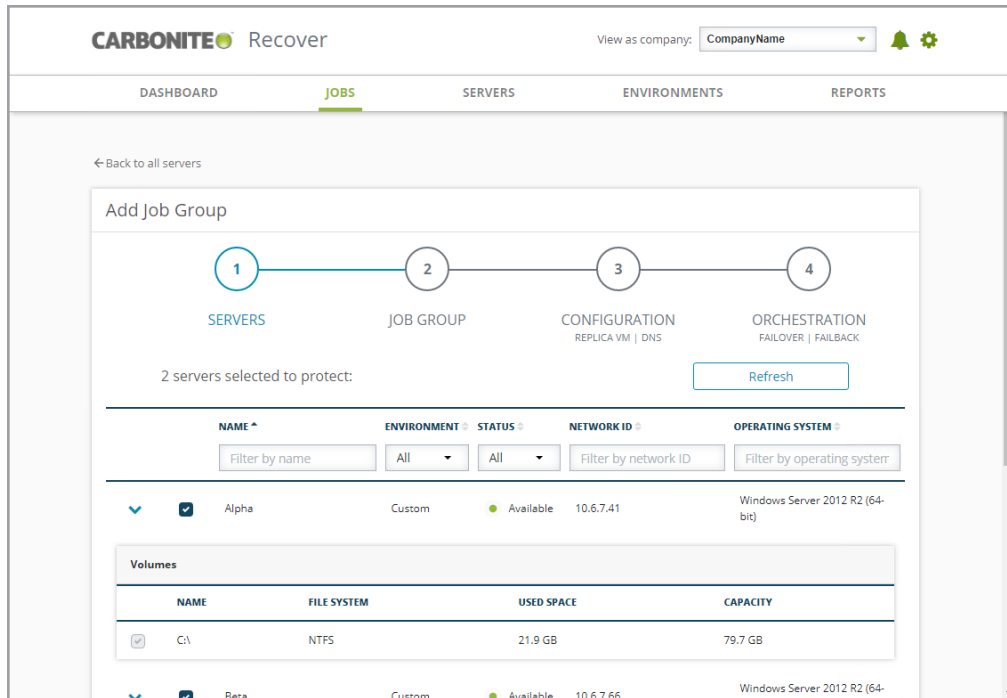
When you create a job, you are creating a job group. That is because you protect servers in groups, even if you have just one server in a group. Groups allow you to manage your servers in orchestration with each other.



Keep the following things in mind when setting up protection.

- If you are protecting domain joined servers, keep in mind the following.
 - You must include one or more domain controllers in your protection group so that the domain servers can resolve DNS and authenticate.
 - The primary domain controller must be authoritative.
 - The primary, authoritative domain controller should be a DNS server.
 - You must configure failover and failback server ordering.
 - The primary, authoritative domain controller must be the first server in the server ordering.
 - All remaining domain controllers must be after the primary, authoritative domain controller and before other domain servers.
 - Make sure you add the replica networking to Windows Sites and Services.
 - If you want to remotely install the replication agent on the servers, you must use domain credentials when adding the server to Carbonite Recover. If you do not use domain credentials, you must install the replication agent on the servers manually.
 - If you select a cloud network that is an isolated network, the servers that are failed over can communicate with each other but may not be able to communicate outside the isolated network. You have a few choices for working with this type of environment.
 - Include all required servers in the job group so all servers failover together.
 - Exclude a server from the job group, but add the required services from that server to another server that is in the job group.
 - Use an SSL VPN client to provide a route to the isolated network. This would be a point-to-site VPN so only the servers running the SSL VPN client will have connectivity.
-

1. To begin the protection process, go to the **Servers** page. Select one or more unprotected source servers and click **Protect** from the overflow menu at the top of the table. To protect just one server, you can also select **Protect** from the overflow menu in a table row for an unprotected source server.
2. On the **Servers** section of the protection wizard, verify the servers you want to protect are selected.

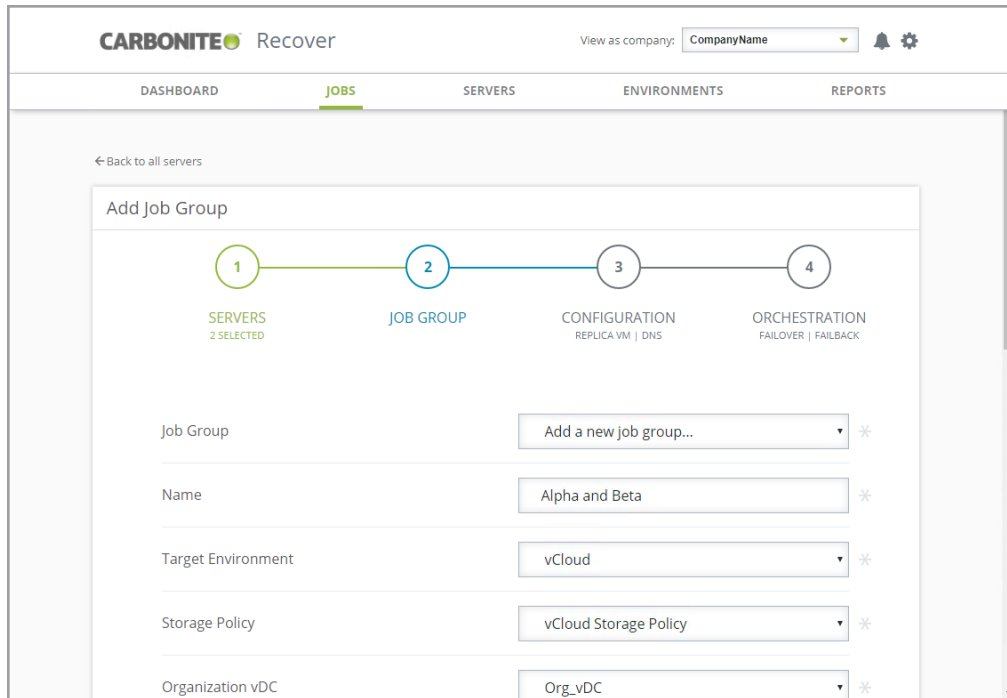


The volumes for each server are listed. You may or may not have options within the **Disks** section, depending on if you have previously protected the server and saved the disks from the previous protection.

- **No option**—If you do not see any options before the volumes, there are no disks available from a previous protection. (The disks must have been saved when the previous protection was deleted and the server must exist within the same environment created in Carbonite Recover in order for previously used disks to be available.)
- **Create new virtual disks**—This option will be available if you have disks available from a previous protection. However, selecting this option will not use these disks. This option will create new disks and the previous disks will be deleted. You may need to use this option if you need to change which disks you are protecting.
- **Reuse existing virtual disks**—This option will be available if you have disks available from a previous protection. This option will reuse the previous disks. Reusing disks makes synchronization faster because only different files need to be synchronized. If you want to reuse disks, you must reuse all of the same disks.

As long as you are not reusing disks, you can exclude data volumes from protection, but be careful when excluding data. Excluded volumes may compromise the integrity of your installed applications. Note the following information about the listed volumes.

- Boot volumes are required for protection and cannot be excluded
 - Unsupported file systems are excluded from protection and cannot be included.
 - The Linux swap disk will be created automatically on the replica, but no data from the swap will be synchronized or replicated.
3. Click **Next** to continue.
 4. On the **Job Group** section of the protection wizard, specify the settings for the job group.



- **Job Group**—Select to create a new job group or choose an existing job group. You can only add to an existing job group if you meet the following criteria.
 - The existing group must be in a healthy, pre-failover state (mirroring or protecting).
 - None of the servers (new to the group or existing in the group) can have pending replication agent upgrades.
 - If the existing group has DNS updates enabled, the servers being added to the group must have valid DNS mappings.

If you are adding to an existing job group, the remaining options on this page of the protection wizard will be read-only.

- **Name**—If you are creating a new group, specify a unique name. If that group name already exists, Carbonite Recover will append a unique number to the group name.
- **Target Environment**—Select the cloud environment where you want to protect the server to.
- **Storage Policy**—Select the storage policy to use from the cloud environment you selected. If you do not know your storage policies, check with Carbonite. This is a default selection for the entire group. You can customize it for each server you are protecting in the next step of the protection wizard.
- **Organization vDC**—Select the organization virtual datacenter to use from the storage policy selected. If you do not know your organization vDC, check with Carbonite. This is a default selection for the entire group. You can customize it for each server you are protecting in the next step of the protection wizard.
- **Windows Appliance**—If you are protecting any Windows servers, select a target appliance to use for those Windows servers. This is a default selection for the entire group. You can customize it for each server you are protecting in the next step of the protection wizard. If you do not have a target appliance, see *Creating a target appliance* on page 28. You will not see this option if you are not protecting any Windows servers.

- **Linux Appliance**—If you are protecting any Linux servers, select a target appliance to use for those Linux servers. This is a default selection for the entire group. You can customize it for each server you are protecting in the next step of the protection wizard. If you do not have a target appliance, see *Creating a target appliance* on page 28. You will not see this option if you are not protecting any Linux servers.
- **Additional Policies**—These policies are group level settings. All jobs in the group will have the same recovery point, bandwidth, and compression settings.
 - **Recovery points**—A recovery point is a snapshot image of the job group taken at a single point in time. You can failover to a recovery point. However, you cannot access the recovery point to recover specific files or folders. Enable this option if you want the ability to take recovery points of the job group on demand. You can also schedule automatic snapshots, if desired. If you disable this option, you will not be able to take any recovery points and will only be able to failover to the most recently replicated data. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.
 - **Maximum to retain**—Specify how many recovery points you want to retain for each job within the group. The minimum is one and the maximum is set by Carbonite. When you reach this limit, Carbonite Recover will delete the oldest recovery point when creating a new one.
 - **Schedule recovery points**—In addition to taking recovery points on demand, enable this option if you want Carbonite Recover to take recovery points of the job group automatically at set intervals. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.
 - **Start**—Specify when you want the recovery point schedule to start.
 - **Repeat every**—Specify how often you want the schedule to take recovery points. The minimum is every one hour.

With Linux, you can take recovery points of system volumes that are non-LVM as long as the root volume is logical, and you can take recovery points of data volumes managed under LVM. Recovery points of LVM data volumes will be created and stored on the target appliance also using LVM.



Recovery points will be skipped until the initial synchronization is complete.

If you enable recovery points and then disable it later, you will be prompted to choose if you want to keep or delete any existing recovery points.

- **Limit bandwidth**—Bandwidth limitations are available to restrict the amount of network bandwidth used for Carbonite Recover data transmissions. When a bandwidth limit is specified, Carbonite Recover never exceeds that allotted amount. The bandwidth not in use by Carbonite Recover is available for all other network traffic. Enable this option if you want to limit bandwidth usage. If you disable this option, Carbonite Recover will use 100% bandwidth availability. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.

- **Limit**—Specify the bandwidth limit to use in Mbps. Carbonite Recover will not exceed that amount. The minimum bandwidth limit is .028 megabits which is 3500 bytes.
- **Schedule bandwidth limit**—Enable this option if you want the **Limit** you have configured to be used only during specific time periods. If you disable this option, the **Limit** will be used all of the time (unless you disable **Limit bandwidth** completely). The toggle circle will be on the right and green when enabled and on the left and gray when disabled.
 - **Start day**—Select the days of the week to apply the bandwidth limit.
 - **Start time**—Select the time of day when you want to start the bandwidth limit.
 - **End time**—Select the time of day when you want to stop the bandwidth limit.



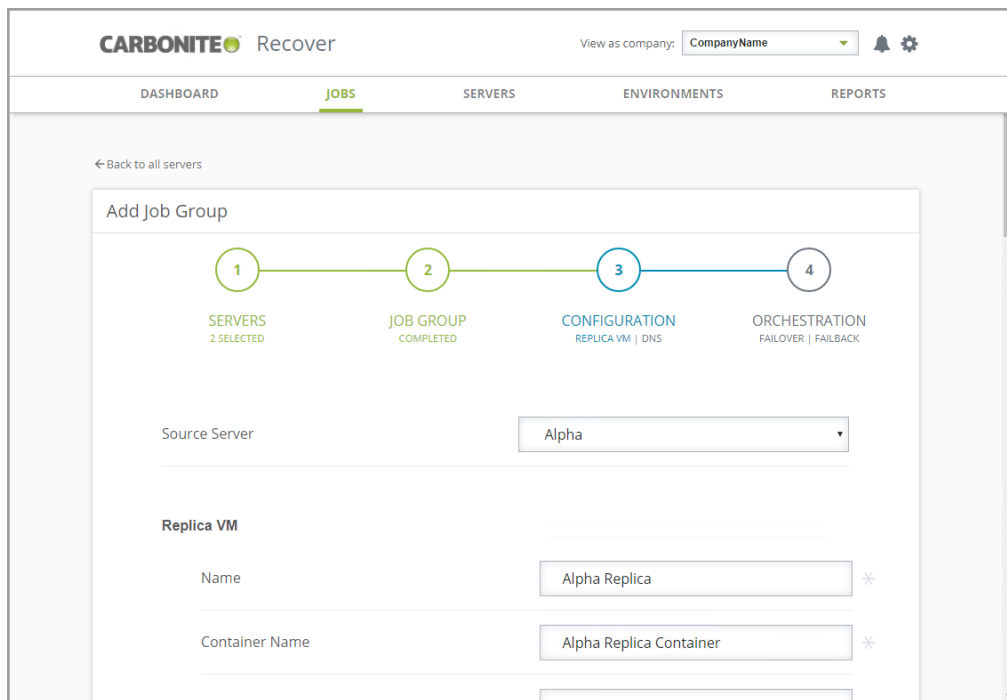
The times specified will be applied using local time on each source server.

If the start time is after the end time, the limit will start at the specified time on the selected days, run overnight, and then stop the next day at the specified time.

- **Compress data**—To help reduce the amount of bandwidth needed to transmit Carbonite Recover data, compression allows you to compress data prior to transmitting it across the network, providing for optimal use of your network resources. The toggle circle will be on the right and green when enabled and on the left and gray when disabled. If compression is enabled, the data is compressed before it is transmitted from the source server. When the target appliance receives the compressed data, it decompresses it and then writes it to disk. Keep in mind that the process of compressing data impacts processor usage on the source server. Use the following guidelines to determine whether you should enable compression.
 - If data is being queued on the source server at any time, consider disabling compression.
 - If the server CPU utilization is averaging over 85%, be cautious about enabling compression.
 - 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 image 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 Recover data.

5. Click **Next** to continue.

6. The **Configuration** section of the protection wizard will vary depending on the operating system of the servers you are protecting.
- **Windows servers only**—If you are only protecting Windows servers, there are two parts to this step of the workflow. You will specify the settings for each individual server you are protecting (using the server name drop-down list at the top of the section) and if you want DNS servers to be updated. **Replica VM** is the first section you will see under the step 3 number.
 - **Linux servers only**—If you are only protecting Linux servers, there is one part to this step of the workflow. You will specify the settings for each individual server you are protecting (using the server name drop-down list at the top of the section). **Replica VM** is the only section you will see under the step 3 number.
 - **Windows and Linux servers**—If you are protecting both Windows and Linux servers, there are two parts to this step of the workflow, however, the second part is only relevant to the Windows servers you are protecting. You will specify the settings for each individual server you are protecting (using the server name drop-down list at the top of the section) and if you want DNS servers to be updated. **Replica VM** is the first section you will see when under the step 3 number.



- **Source Server**—Select a server from the list and then configure the options on the rest of the page for that selected server. Repeat this process for each server in the list. If you are adding to an existing job group, the options for the existing servers will be read-only.
- **Replica VM**—Specify how you want the replica virtual machine to be created in the cloud during failover.
 - **Name**—Specify the virtual machine display name. This is the replica virtual machine that will be created in the cloud. By default, this is the name of the original source server with the suffix Replica.

- **Container Name**—Specify the name of the container to create in the cloud. If the name already exists, Carbonite Recover will append a unique number to the name.
- **Size**—Select the size of the replica virtual machine. You can select **Specify** and identify the amount of **Memory** and the number of **Cores/sockets** for the replica virtual machine. You may also have predefined sizes set by Carbonite. If you have predefined sizes but are uncertain what the specifications are for the size, contact Carbonite
- **Storage**—Specify how you want to handle storage on the replica virtual machine.
 - **Target Appliance**—The target appliance you selected on the previous page of the protection wizard will be selected. If desired, you can select a different target appliance for an individual server.
 - **Virtual Disks**—For each volume on the source server, specify how large you want the corresponding volume to be on the replica virtual machine. The replica disk size cannot be smaller than the total disk size on the source volume. It must be as large or larger than the total source disk size. If you are reusing disks from a previous protection, you cannot modify the disk selection or disk size.

If you use the default replica disk size when creating a job, a moderate amount of extra disk capacity is automatically added to accommodate recovery points. If you customize the disk size, no additional disk capacity is added. If you expect a high rate of data change, customize the disk size to increase the capacity for larger recovery points.

- **Exclusions**—For the volumes you have protected, you can enter exclusions for data you do not want protected. For example, if you have a temporary database on your protected volume but do not need to protect it, you can exclude it here.
 - **Path**—Enter a path, specific file, or a wildcard to be excluded from protection.
 - **Recursive**—Select this option if you want the exclusion to be recursive, which indicates the exclusion will be automatically applied to the sub-directories of the specified path.
 - **Add**—Click this button to save the entered exclusion rule.
 - **Delete**—Click the trash icon to remove a saved exclusion rule.
- **Network**—Networking is separated into two sections. The **Failover Network** section contains settings that will be used for live and recovery point failover. The **Test Failover Network** contains settings that will be used for test failover. For each type of failover, specify the networking values to be used on the replica virtual machine, and for each adapter on the source server, specify how you want the corresponding adapter to be configured on the replica virtual machine.
 - **Network**—Select the network that you want the adapter to use on the replica virtual machine. Keep in mind the following when selecting a network.
 - If you have not been assigned two networks or you do not understand your available networks, check with Carbonite.
 - If you select a cloud network that is an isolated network, the servers that are failed over can communicate with each other but may not be able to

communicate outside the isolated network. You have a few choices for working with this type of environment.

- Include all required servers in the job group so all servers failover together.
 - Exclude a server from the job group, but add the required services from that server to another server that is in the job group.
 - Use an SSL VPN client to provide a route to the isolated network. This would be a point-to-site VPN so only the servers running the SSL VPN client will have connectivity.
- **Subnet**—The subnet should default to the correct setting based on the network you selected. If needed, specify a different subnet to assign to the adapters. Make sure the subnet is valid for the selected **Network**.
 - **Gateway**—The gateway should default to the correct setting based on the network you selected. If needed, specify a different gateway to assign to the adapters. Make sure the gateway is valid for the selected **Network**.
 - **DNS Addresses**—The DNS addresses should default to the correct settings based on the network you selected. If needed, specify different DNS addresses to apply to the adapter. Make sure the DNS addresses are valid for the selected **Network** and your replica virtual machines will be able to reach the DNS server after failover. Click **+** to add another row to the table or **-** to remove an existing row from the table. Specify them in the order you want them used.
 - **Adapter Type**—For each adapter on the source, select the network adapter type you want to use on the replica. The types available in the list will depend on the operating system you have selected, and if you have appropriate utilities, such as VMware Tools, installed on your source server.
 - **IP Mode**—For each adapter on the source, select how you want the IP address to be assigned on the replica. Select **Pool** if you want the replica virtual machine to be assigned an IP address from a pool of addresses. Select **Manual** and then specify an IP address, if you want to assign a specific IP address to the adapter. The specified address must be valid for the selected **Network** and it cannot be allocated already. Click **+** to add another row to the table or **-** to remove an existing row from the table.

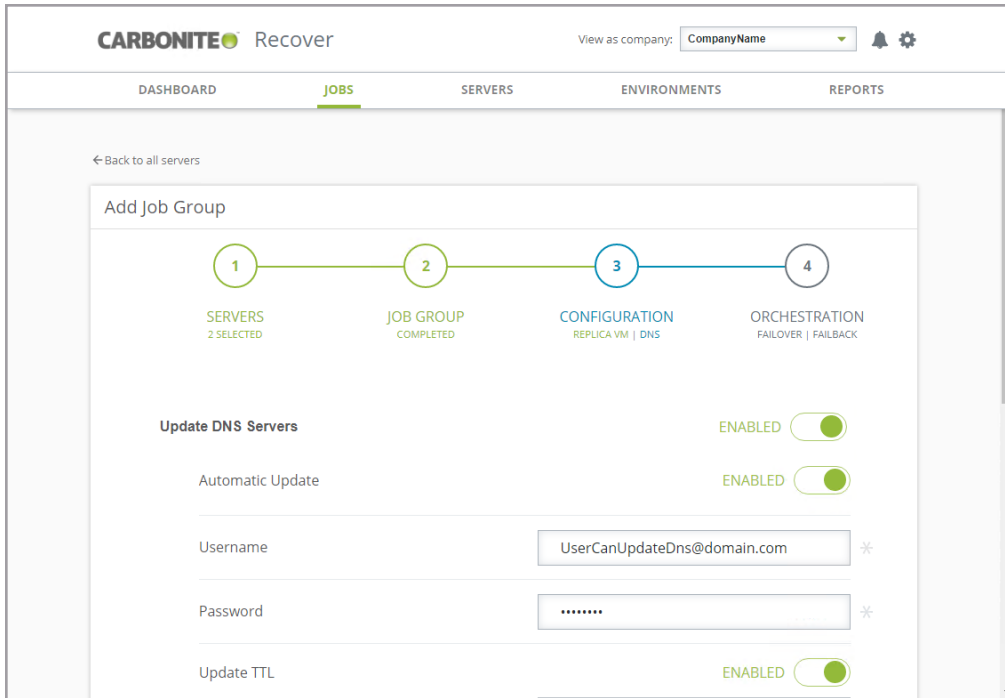


You should be protecting a source DNS server in order to be able to access other protected source servers after failover. Carbonite will provide you with an IP address to specify for your replica DNS server in the cloud. If you do not know the IP address to use, check with Carbonite.

Network updates made during failover will be based on the network adapter name when the job 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.

7. Click **Next** to continue.
8. If you are protecting Windows servers only or a mix of Windows and Linux servers, you will now be at the **DNS** section under the step 3 number. If you are protecting only Linux servers, you will

not have this section. Also, if you are adding to an existing job group, the DNS settings will be read-only.



- **Update DNS Servers**—Enable this option if you want Carbonite Recover to update your DNS servers during failover. The toggle circle will be on the right and green when enabled and on the left and gray when disabled. This option allows you to update DNS automatically at failover time or you can trigger the updates manually. If you disable this option, you will not be able to update DNS automatically or manually.
- **Automatic Update**—Enable this option if you want the DNS updates to happen automatically during failover. The toggle circle will be on the right and green when enabled and on the left and gray when disabled. Whether automatic updates are enabled or disabled, you can still trigger DNSU updates manually.
- **Username**—Specify a user, in the format name@domain.com, 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 on the source's A (host) and PTR (reverse lookup) records. These permissions are not included by default in the DnsAdmins group.
- **Password**—Specify the password associated with the user.
- **Update TTL**—When enabled, Carbonite Recover will update the time to live when DNS updates are made. The toggle circle will be on the right and green when enabled and on the left and gray when disabled. Specify the length of time, in seconds, for the TTL value for all modified DNS A records. Ideally, you should specify 300 seconds (5 minutes) or less.
- **DNS Servers**—The list of DNS servers is populated from the DNS servers associated with the source servers you are protecting and the failover networks in the cloud that you have selected for each NIC on those source servers (step 3 individual server settings). If

you do not want to update one of these DNS servers, remove it from the list by clicking the minus icon. If you want to add a DNS server that is not in the list, click the plus icon and enter the IP address.

- **DNS Entries**— If the replica IP address after failover was configured for pool, then choose **Auto** to automatically use the assigned pool address after failover. If the replica IP address after failover was configured manually for a specific address, select the address you want DNS to use after failover. In either case, you can also set a source IP address to **Discard**.
9. Click **Next** to continue.
 10. On the **Orchestration** section of the protection wizard, specify the failover, restore, and failback settings for the group. There are two parts to this step of the workflow. **Failover** is the first section under the step 4 number.



The fields on this page will vary depending on if your job group contains one server or more than one server.

If you are adding to an existing job group, the orchestration settings will be read-only. The new servers will be added to the end of the server ordering. You can make modifications after the new servers have been added.

The screenshot shows the CARBONITE Recover web interface. At the top, there's a navigation bar with 'DASHBOARD', 'JOBS', 'SERVERS', 'ENVIRONMENTS', and 'REPORTS'. The 'JOBS' tab is active. Below the navigation bar, there's a 'Back to all servers' link. The main content area is titled 'Add Job Group' and shows a progress bar with four steps: 1. SERVERS (2 SELECTED), 2. JOB GROUP (COMPLETED), 3. CONFIGURATION (REPLICA VM | DNS), and 4. ORCHESTRATION (FAILOVER | FAILBACK). Below the progress bar, there's a section for 'Pre-Failover Script' with a 'Delete script' button. The 'Script name' field contains 'PreFailoverScript.ps1' and has a 'Browse' button. The 'Arguments' field contains '-FirstParameter 'value1' -SecondParamet'. The 'Description' field contains 'Script to execute task1 and task2'.

- **Pre-Failover Script**—Before failover starts, you can have your own script launched on a particular server. This script must be a PowerShell or Linux bash script.
 - **Script name**—Browse (by default the local machine) and select the script that you want to run before the failover process starts. Once you select a script, the rest of

the script fields will be displayed. The selected script will be encrypted and uploaded to be executed on the specified server at the designated time.

- **Arguments**—Identify any arguments that you want passed into the script. If you are using a PowerShell script, the argument list must be a space separated list and follow standard PowerShell usage rules for hyphens, string quoting, and using special characters. For example, your PowerShell arguments might be -parameter1 'value1' -parameter2 'value2'. If you are using a bash script, the argument list must be a space separated list and follow standard bash usage rules for string quoting and using special characters. For example, your bash arguments might be "value1 'value2' value3".
- **Description**—You must add a unique description to the script. The description is used to identify the script.
- **Run script on**—Select the server where you want the script to run. Keep in mind, workers cannot execute bash scripts. They can only execute PowerShell scripts. Also, if you select a source server and that server is down, the pre-failover script will not be able to be run.
- **If script fails, continue with failover**—If a script does not complete within ten minutes, the script will be considered a failure. Additionally, if there are any failures while the script is executing, the script will be considered a failure. The failover process can continue even if the script execution fails. If you disable this option, a script failure will stop the failover process. You will have to fix the script failure and restart the failover process. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.



Click **Delete script** to remove a script you have already specified.

You will have the opportunity to disable or change scripts before the failover process is started, if desired.

- **Failover Order**—If you have more than one server in your group, you can set the failover order of the servers. You will not see this section if your job group contains only one server.
 - **Use startup order for failover and fallback**—You will only see this option if you are not using scripts. Enable this option to allow for server ordering. The toggle circle will be on the right and green when enabled and on the left and gray when disabled. If you are using scripts, you will not see this option because server ordering is automatically enabled.

Drag and drop the servers in the group to the order you want them failed over. If you are using scripts, the pre-failover script will be associated with the first server in the list and the post-failover script will be associated with the last server in the list. (Associated meaning the script is executed when it is that server's turn in the server order, not that the script will run on that server.) Servers in the list will not power on until the replica virtual machine before it in the startup order has completed all of its failover process operations.



If you are protecting domain joined servers, keep in mind the following.

- You must include one or more domain controllers in your protection group so that the domain servers can resolve DNS and authenticate.
- The primary domain controller must be authoritative.
- The primary, authoritative domain controller should be a DNS server.
- You must configure failover and failback server ordering.
- The primary, authoritative domain controller must be the first server in the server ordering.
- All remaining domain controllers must be after the primary, authoritative domain controller and before other domain servers.
- Make sure you add the replica networking to Windows Sites and Services.
- If you want to remotely install the replication agent on the servers, you must use domain credentials when adding the server to Carbonite Recover. If you do not use domain credentials, you must install the replication agent on the servers manually.

You will have the opportunity to rearrange or disable the startup order before the failover process is started, if desired.

- **Post-Failover Script**—After the failover process is completed (when the last server in the startup order is online), you can have your own script launched on a particular server. This script must be a PowerShell or Linux bash script.
 - **Script Name**—Browse (by default the local machine) and select the script that you want to run after the failover process completes. Once you select a script, the rest of the script fields will be displayed. The selected script will be encrypted and uploaded to be executed on the specified server at the designated time.
 - **Arguments**—Identify any arguments that you want passed into the script. If you are using a PowerShell script, the argument list must be a space separated list and follow standard PowerShell usage rules for hyphens, string quoting, and using special characters. For example, your PowerShell arguments might be -parameter1 'value1' -parameter2 'value2'. If you are using a bash script, the argument list must be a space separated list and follow standard bash usage rules for string quoting and using special characters. For example, your bash arguments might be "value1 'value2' value3".
 - **Description**—You must add a unique description to the script. The description is used to identify the script.
 - **Run Script On**—Select the server where you want the script to run. Keep in mind, workers cannot execute bash scripts. They can only execute PowerShell scripts.



Click **Delete script** if you need to remove a script you have already specified.

You will have the opportunity to disable or change scripts before the failover process is started, if desired.

11. Click **Next** to continue.
12. On the **Orchestration** section of the protection wizard, specify the failover, restore, and failback settings for the group. There are two parts to this step of the workflow. **Failback** is the second section under the step 4 number.



The fields on this page will vary depending on if your job group contains one server or more than one server.

If you are adding to an existing job group, the orchestration settings will be read-only. The new servers will be added to the end of the server ordering. You can make modifications after the new servers have been added.

The screenshot shows the CARBONITE Recover interface. At the top, there's a navigation bar with 'DASHBOARD', 'JOBS', 'SERVERS', 'ENVIRONMENTS', and 'REPORTS'. The 'JOBS' tab is active. Below the navigation bar, there's a 'Back to all servers' link. The main content area is titled 'Add Job Group' and shows a progress bar with four steps: 1. SERVERS (2 SELECTED), 2. JOB GROUP (COMPLETED), 3. CONFIGURATION (REPLICA VM | DNS), and 4. ORCHESTRATION (FAILOVER | FAILBACK). Below the progress bar, there's a section for adding pre- and post-failback scripts. The 'Pre-Failback Script' section has a 'Delete script' button and a 'Browse' button. The 'Script name' field contains 'PreFailbackScript.ps1'. The 'Arguments' field contains '-FirstParameter 'value1''. The 'Description' field contains 'Script to execute task1'.

- **Pre-Failback Script**—Before failback starts, you can have your own script launched on a particular server. This script must be a PowerShell or Linux bash script.
 - **Script Name**—Browse (by default the local machine) and select the script that you want to run before the failback process starts. Once you select a script, the rest of the script fields will be displayed. The selected script will be encrypted and uploaded to be executed on the specified server at the designated time.
 - **Arguments**—Identify any arguments that you want passed into the script. If you are using a PowerShell script, the argument list must be a space separated list and

follow standard PowerShell usage rules for hyphens, string quoting, and using special characters. For example, your PowerShell arguments might be -parameter1 'value1' -parameter2 'value2'. If you are using a bash script, the argument list must be a space separated list and follow standard bash usage rules for string quoting and using special characters. For example, your bash arguments might be "value1 'value2' value3".

- **Description**—You must add a unique description to the script. The description is used to identify the script.
- **Run Script On**—Select the server where you want the script to run. Keep in mind, workers cannot execute bash scripts. They can only execute PowerShell scripts.
- **If script fails, continue with restore**—If a script does not complete within ten minutes, the script will be considered a failure. Additionally, if there are any failures while the script is executing, the script will be considered a failure. The failback process can continue even if the script execution fails. If you disable this option, a script failure will stop the failback process. You will have to fix the script failure and restart the failback process. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.



Click **Delete Script** if you need to remove a script you have already specified.

You will have the opportunity to disable or change scripts before the restore process is started, if desired.

- **Failback Order**—Your server failback order will match the failover order you configured. You will have the opportunity to rearrange or disable the failback startup order before the restore process is started, if desired. Servers in the list will not start the failback process until the server before it in the startup order has completed all of its failback process operations.
- **Post-Failback Script**—After the failback process is completed (when the last server in the startup order is online), you can have your own script launched on a particular server. This script must be a PowerShell or Linux bash script.
 - **Script Name**—Browse (by default the local machine) and select the script that you want to run after the failback process completes. Once you select a script, the rest of the script fields will be displayed. The selected script will be encrypted and uploaded to be executed on the specified server at the designated time.
 - **Arguments**—Identify any arguments that you want passed into the script. If you are using a PowerShell script, the argument list must be a space separated list and follow standard PowerShell usage rules for hyphens, string quoting, and using special characters. For example, your PowerShell arguments might be -parameter1 'value1' -parameter2 'value2'. If you are using a bash script, the argument list must be a space separated list and follow standard bash usage rules for string quoting and using special characters. For example, your bash arguments might be "value1 'value2' value3".

- **Description**—You must add a unique description to the script. The description is used to identify the script.
- **Run Script On**—Select the server where you want the script to run. Keep in mind, workers cannot execute bash scripts. They can only execute PowerShell scripts.



Click **Delete script** if you need to remove a script you have already specified.

You will have the opportunity to disable or change scripts before the restore process is started, if desired.

13. Click **Verify** to continue.
14. Carbonite Recover validates settings for each target appliance and source server. The **Verification Checklist** page displays the validation items. Expand a target appliance or source server name to see the validation items associated with that server.

Errors are designated by a white X inside a red circle. Warnings are designated by a white exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle.

Depending on the warning or error, you may see a button allowing you to **Fix** or **Fix All**. This will allow Carbonite Recover to correct the problem for you. For those warnings or errors that Carbonite Recover cannot correct automatically or any fixes that could not be successfully completed, you will need to manually correct the problem. You can revalidate the servers by clicking **Recheck**.

You can also search for specific items by using the filter.

You can continue with warnings, however, you must correct any errors before you can continue.

15. Once your configuration has passed verification with no errors, click **Finish** to start protection.



If you are protecting a domain controller or DNS server, make sure you have added the target networking to Windows Sites and Services.

Testing failover

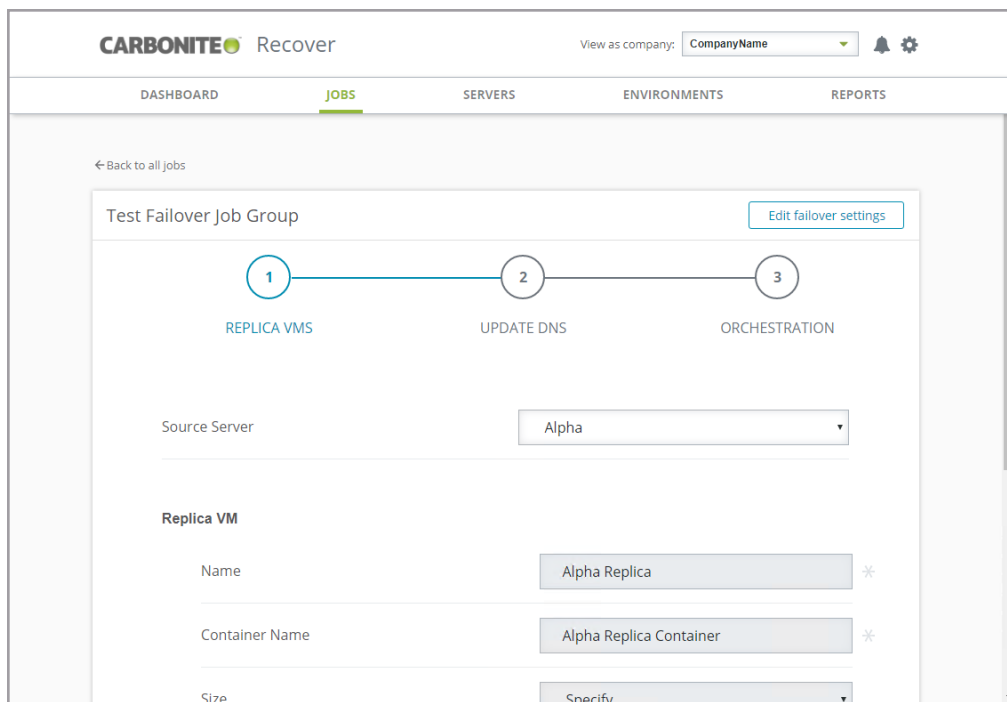
In the event you need to test failover, you can quickly perform a test to the cloud. The test failover process workflow is different if you are testing at the group level or at an individual level. The workflow also varies if you are failing over only Linux servers.

- **Group level**—Use the group level instructions if you are failing over a group of multiple servers or if you are failing over a group that has only one server.
- **Individual level**—Use the individual instructions if you are failing over an individual server from a group of multiple servers.

Testing failover at the group level

Use this process if you are testing failover for a group of multiple servers or if you are testing failover for a group that has only one server. If you are testing failover for an individual server from a group of multiple servers, use the testing failover at the individual level instructions.

1. On the **Jobs** page, select **Test Failover** from the overflow menu for the group. Group actions will only be available when all servers in the group can safely perform that action. If you have only one server in a group, you will only have group actions.
2. You will see the failover options you specified when you created the job. If you want to edit the options, click **Edit failover settings**.



- **Source Server**—Select a server from the list and then configure the options on the rest of the page for that selected server. Repeat this process for each server in the list.
- **Authoritative failover**—This option will only be displayed if you are failing over a backup domain controller that is running in functional level 2008 R2 or older and your failover group does not contain a primary domain controller. In this specific case, you can enable

this option so that the backup domain controller will be authoritative after failover. However, you should use this option with caution. Enable it when performing a live or recovery point failover and you need to rebuild your domain in the cloud. You can also enable it when performing a test failover to an isolated network. However, disable this option when your primary domain controller is still online or you will corrupt your source domain. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.

- **Replica VM**—Specify how you want the replica virtual machine to be created in the cloud during failover.
 - **Name**—Specify the virtual machine display name. This is the replica virtual machine that will be created in the cloud. By default, this is the name of the original source server with the suffix Replica.
 - **Container Name**—Specify the name of the container to create in the cloud. If the name already exists, Carbonite Recover will append a unique number to the name.
 - **Size**—Select the size of the replica virtual machine. You can select **Specify** and identify the amount of **Memory** and the number of **Cores/Sockets** for the replica virtual machine. You may also have predefined sizes set by Carbonite. If you have predefined sizes but are uncertain what the specifications are for the size, contact Carbonite
- **Test Failover Network**—These are the networking values to be used on the replica virtual machine during the test failover.
 - **Network**—Select the network that you want the adapter to use on the replica virtual machine. Keep in mind the following when selecting a network.
 - If you have not been assigned two networks or you do not understand your available networks, check with Carbonite.
 - If you select a cloud network that is an isolated network, the servers that are failed over can communicate with each other but may not be able to communicate outside the isolated network. You have a few choices for working with this type of environment.
 - Include all required servers in the job group so all servers failover together.
 - Exclude a server from the job group, but add the required services from that server to another server that is in the job group.
 - Use an SSL VPN client to provide a route to the isolated network. This would be a point-to-site VPN so only the servers running the SSL VPN client will have connectivity.
 - **Subnet**—The subnet should default to the correct setting based on the network you selected. If needed, specify a different subnet to assign to the adapters. Make sure the subnet is valid for the selected **Network**.
 - **Gateway**—The gateway should default to the correct setting based on the network you selected. If needed, specify a different gateway to assign to the adapters. Make sure the gateway is valid for the selected **Network**.
 - **DNS Addresses**—The DNS addresses should default to the correct settings based on the network you selected. If needed, specify different DNS addresses to

apply to the adapter. Make sure the DNS addresses are valid for the selected **Network** and your replica virtual machines will be able to reach the DNS server after failover. Click **+** to add another row to the table or **-** to remove an existing row from the table. Specify them in the order you want them used.

- **Adapter Type**—For each adapter on the source, select the network adapter type you want to use on the replica. The types available in the list will depend on the operating system you have selected, and if you have appropriate utilities, such as VMware Tools, installed on your source server.
- **IP Mode**—For each adapter on the source, select how you want the IP address to be assigned on the replica. Select **Pool** if you want the replica virtual machine to be assigned an IP address from a pool of addresses. Select **Manual** and then specify an IP address, if you want to assign a specific IP address to the adapter. The specified address must be valid for the selected **Network** and it cannot be allocated already. Click **+** to add another row to the table or **-** to remove an existing row from the table.



You should be protecting a source DNS server in order to be able to access other protected source servers after failover. Carbonite will provide you with an IP address to specify for your replica DNS server in the cloud. If you do not know the IP address to use, check with Carbonite.

3. Click **Next** to continue.
4. DNS updates are not available for test failover. Click **Next** to continue.
5. Review your orchestration settings. If you need to modify them, click **Edit failover settings**.

The screenshot shows the Carbonite Recover web interface. At the top, there is a navigation bar with 'CARBONITE Recover' and a 'View as company:' dropdown menu. Below the navigation bar, there are tabs for 'DASHBOARD', 'JOBS', 'SERVERS', 'ENVIRONMENTS', and 'REPORTS'. The 'JOBS' tab is selected. The main content area shows a 'Test Failover Job Group' card. The card has a progress flow with three steps: 1. REPLICAS VMS COMPLETED, 2. UPDATE DNS COMPLETED, and 3. ORCHESTRATION. Below the flow, there are settings for 'Use failover plan' (ENABLED) and 'Pre-Failover Script' (ENABLED). The description for the Pre-Failover Script is 'Script to execute task1 and task2'. An 'Edit failover plan' button is visible in the top right corner of the job group card.

- **Use failover plan**—Enable this option to start the servers in the order specified. You have the option of enabling or disabling scripts as desired. If you disable this option, the servers will all start at the same time and the scripts will be automatically disabled. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.
- **Edit plan**—This link allows you to edit any of your existing plan settings. You can change the script settings or server order.
 - **Pre-Failover Script**—Before failover starts, you can have your own script launched on a particular server. This script must be a PowerShell or Linux bash script. If you are using the failover plan, you can disable or enable scripts so they do or do not run. If you are not using the failover plan, scripts will automatically be disabled. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.
 - **Script name**—Browse (by default the local machine) and select the script that you want to run before the failover process starts. Once you select a script, the rest of the script fields will be displayed. The selected script will be encrypted and uploaded to be executed on the specified server at the designated time.
 - **Arguments**—Identify any arguments that you want passed into the script. If you are using a PowerShell script, the argument list must be a space separated list and follow standard PowerShell usage rules for hyphens, string quoting, and using special characters. For example, your PowerShell arguments might be `-parameter1 'value1' -parameter2 'value2'`. If you are using a bash script, the argument list must be a space separated list and follow standard bash usage rules for string quoting and using special characters. For example, your bash arguments might be `"value1 'value2' value3"`.
 - **Description**—You must add a unique description to the script. The description is used to identify the script.
 - **Run script on**—Select the server where you want the script to run. Keep in mind, workers cannot execute bash scripts. They can only execute PowerShell scripts. Also, if you select a source server and that server is down, the pre-failover script will not be able to be run.
 - **If script fails, continue with failover**—If a script does not complete within ten minutes, the script will be considered a failure. Additionally, if there are any failures while the script is executing, the script will be considered a failure. The failover process can continue even if the script execution fails. If you disable this option, a script failure will stop the failover process. You will have to fix the script failure and restart the failover process. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.



Click **Delete script** to remove a script you have already specified.

- **Failover Order**— You can set the failover order of the servers. You will not see this section if your job group contains only one server.
 - **Use startup order for failover and fallback**—You will only see this option if you are not using scripts. Enable this option to allow for server ordering. The toggle circle will be on the right and green when enabled and on the left and gray when disabled. If you are using scripts, you will not see this option because server ordering is automatically enabled.

Drag and drop the servers in the group to the order you want them failed over. If you are using scripts, the pre-failover script will be associated with the first server in the list and the post-failover script will be associated with the last server in the list. (Associated meaning the script is executed when it is that server's turn in the server order, not that the script will run on that server.) Servers in the list will not power on until the replica virtual machine before it in the startup order has completed all of its failover process operations.



If you are protecting domain joined servers, keep in mind the following.

- You must include one or more domain controllers in your protection group so that the domain servers can resolve DNS and authenticate.
- The primary domain controller must be authoritative.
- The primary, authoritative domain controller should be a DNS server.
- You must configure failover and fallback server ordering.
- The primary, authoritative domain controller must be the first server in the server ordering.
- All remaining domain controllers must be after the primary, authoritative domain controller and before other domain servers.
- Make sure you add the replica networking to Windows Sites and Services.
- If you want to remotely install the replication agent on the servers, you must use domain credentials when adding the server to Carbonite Recover. If you do not use domain credentials, you must install the replication agent on the servers manually.

-
- **Post-Failover Script**—After the failover process is completed (when the last server in the startup order is online), you can have your own script launched on a particular server. This script must be a PowerShell or Linux bash script. If you are using the failover plan, you can disable or enable scripts so they do or do not run. If you are not using the failover plan, scripts will automatically be disabled. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.

- **Script Name**—Browse (by default the local machine) and select the script that you want to run after the failover process completes. Once you select a script, the rest of the script fields will be displayed. The selected script will be encrypted and uploaded to be executed on the specified server at the designated time.
- **Arguments**—Identify any arguments that you want passed into the script. If you are using a PowerShell script, the argument list must be a space separated list and follow standard PowerShell usage rules for hyphens, string quoting, and using special characters. For example, your PowerShell arguments might be `-parameter1 'value1' -parameter2 'value2'`. If you are using a bash script, the argument list must be a space separated list and follow standard bash usage rules for string quoting and using special characters. For example, your bash arguments might be `"value1 'value2' value3"`.
- **Description**—You must add a unique description to the script. The description is used to identify the script.
- **Run Script On**—Select the server where you want the script to run. Keep in mind, workers cannot execute bash scripts. They can only execute PowerShell scripts.



Click **Delete script** if you need to remove a script you have already specified.

6. Click **Verify** to continue.
7. Carbonite Recover validates settings for each source server. The **Verification Checklist** page displays the validation items. Expand a source server name to see the validation items associated with that server.

Errors are designated by a white X inside a red circle. Warnings are designated by a white exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle.

Depending on the warning or error, you may see a button allowing you to **Fix** or **Fix All**. This will allow Carbonite Recover to correct the problem for you. For those warnings or errors that Carbonite Recover cannot correct automatically or any fixes that could not be successfully completed, you will need to manually correct the problem. You can revalidate the servers by clicking **Recheck**.

You can also search for specific items by using the filter.

You can continue with warnings, however, you must correct any errors before you can continue.

8. Once your configuration has passed verification with no errors, click **Finish** to start the failover.



If you are failing over to an isolated network and your target environment worker does not have access to that network, Carbonite Recover will not be able to verify the readiness of the replica virtual machine or complete other post-failover tasks.

Testing failover at the individual level

Use this process if you are testing failover for an individual server from a group of multiple servers. If you are testing failover for a group of multiple servers or if you are testing failover for a group that has only one server, use the testing failover at the group level instructions.

1. On the **Jobs** page, select **Test Failover** from the overflow menu for an individual server.
2. You will see the failover options you specified when you created the job. If you want to edit the options, click **Edit failover settings**.



If you are using scripts, the pre-failover script will be associated with the first server in the list and the post-failover script will be associated with the last server in the list. (Associated meaning the script is executed when it is that server's turn in the server order, not that the script will run on that server.) If you are failing over either the first or last server by itself, you will see a warning on the **Test Failover Job** page. It is a notification that scripts will not be run since you are only failing over that single server and not the group.

- **Replica VM**—Specify how you want the replica virtual machine to be created in the cloud during failover.
 - **Name**—Specify the virtual machine display name. This is the replica virtual machine that will be created in the cloud. By default, this is the name of the original source server with the suffix Replica.

- **Container Name**—Specify the name of the container to create in the cloud. If the name already exists, Carbonite Recover will append a unique number to the name.
- **Size**—Select the size of the replica virtual machine. You can select **Specify** and identify the amount of **Memory** and the number of **Cores/Sockets** for the replica virtual machine. You may also have predefined sizes set by Carbonite. If you have predefined sizes but are uncertain what the specifications are for the size, contact Carbonite
- **Authoritative failover**—This option will only be displayed if you are failing over a backup domain controller that is running in functional level 2008 R2 or older and your failover group does not contain a primary domain controller. In this specific case, you can enable this option so that the backup domain controller will be authoritative after failover. However, you should use this option with caution. Enable it when performing a live or recovery point failover and you need to rebuild your domain in the cloud. You can also enable it when performing a test failover to an isolated network. However, disable this option when your primary domain controller is still online or you will corrupt your source domain. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.
- **Test Failover Network**—These are the networking values to be used on the replica virtual machine during the test failover.
 - **Network**—Select the network that you want the adapter to use on the replica virtual machine. Keep in mind the following when selecting a network.
 - If you have not been assigned two networks or you do not understand your available networks, check with Carbonite.
 - If you select a cloud network that is an isolated network, the servers that are failed over can communicate with each other but may not be able to communicate outside the isolated network. You have a few choices for working with this type of environment.
 - Include all required servers in the job group so all servers failover together.
 - Exclude a server from the job group, but add the required services from that server to another server that is in the job group.
 - Use an SSL VPN client to provide a route to the isolated network. This would be a point-to-site VPN so only the servers running the SSL VPN client will have connectivity.
 - **Subnet**—The subnet should default to the correct setting based on the network you selected. If needed, specify a different subnet to assign to the adapters. Make sure the subnet is valid for the selected **Network**.
 - **Gateway**—The gateway should default to the correct setting based on the network you selected. If needed, specify a different gateway to assign to the adapters. Make sure the gateway is valid for the selected **Network**.
 - **DNS Addresses**—The DNS addresses should default to the correct settings based on the network you selected. If needed, specify different DNS addresses to apply to the adapter. Make sure the DNS addresses are valid for the selected **Network** and your replica virtual machines will be able to reach the DNS server

after failover. Click **+** to add another row to the table or **-** to remove an existing row from the table. Specify them in the order you want them used.

- **Adapter Type**—For each adapter on the source, select the network adapter type you want to use on the replica. The types available in the list will depend on the operating system you have selected, and if you have appropriate utilities, such as VMware Tools, installed on your source server.
- **IP Mode**—For each adapter on the source, select how you want the IP address to be assigned on the replica. Select **Pool** if you want the replica virtual machine to be assigned an IP address from a pool of addresses. Select **Manual** and then specify an IP address, if you want to assign a specific IP address to the adapter. The specified address must be valid for the selected **Network** and it cannot be allocated already. Click **+** to add another row to the table or **-** to remove an existing row from the table.



You should be protecting a source DNS server in order to be able to access other protected source servers after failover. Carbonite will provide you with an IP address to specify for your replica DNS server in the cloud. If you do not know the IP address to use, check with Carbonite.

3. Click **Next** to continue.
4. DNS updates are not available for test failover. Click **Verify** to continue.
5. Carbonite Recover validates settings for each source server. The **Verification Checklist** page displays the validation items. Expand a source server name to see the validation items associated with that server.

Errors are designated by a white X inside a red circle. Warnings are designated by a white exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle.

Depending on the warning or error, you may see a button allowing you to **Fix** or **Fix All**. This will allow Carbonite Recover to correct the problem for you. For those warnings or errors that Carbonite Recover cannot correct automatically or any fixes that could not be successfully completed, you will need to manually correct the problem. You can revalidate the servers by clicking **Recheck**.

You can also search for specific items by using the filter.

You can continue with warnings, however, you must correct any errors before you can continue.

6. Once your configuration has passed verification with no errors, click **Finish** to start the test failover.



If you are failing over to an isolated network and your target environment worker does not have access to that network, Carbonite Recover will not be able to verify the readiness of the replica virtual machine or complete other post-failover tasks.

Failing over

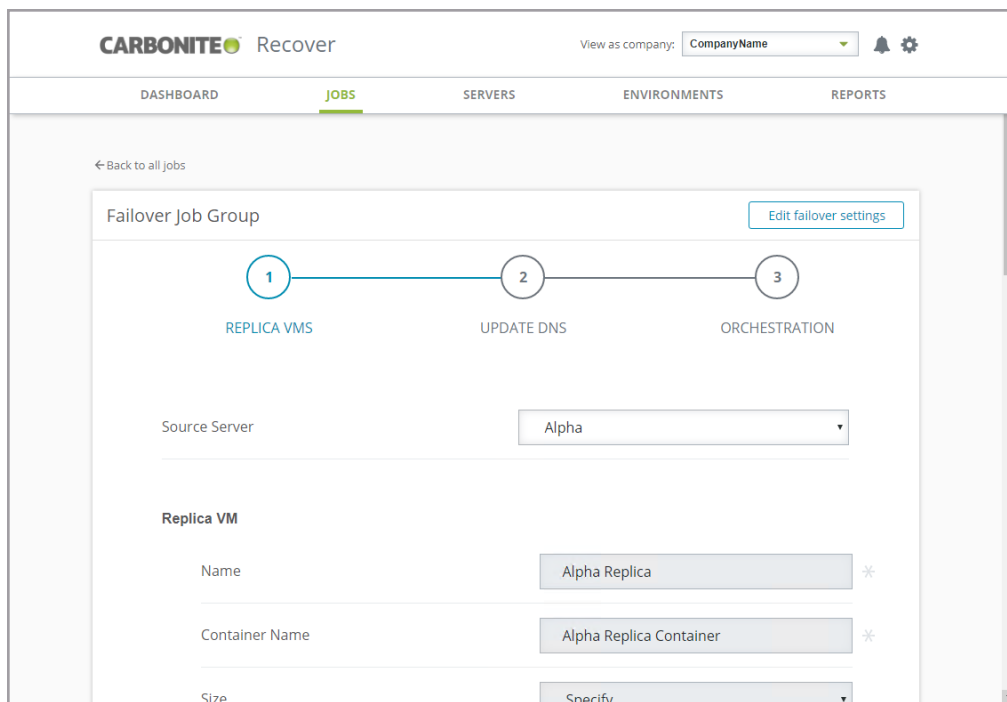
You can quickly fail over one or more source servers in the cloud. The failover process workflow is different if you are failing over at the group level or at an individual level. The workflow also varies if you are failing over only Linux servers.

- **Group level**—Use the group level instructions if you are failing over a group of multiple servers or if you are failing over a group that has only one server.
- **Individual level**—Use the individual instructions if you are failing over an individual server from a group of multiple servers.

Failing over at the group level

Use this process if you are failing over a group of multiple servers or if you are failing over a group with only one server. If you are failing over an individual server from a group of multiple servers, use the failing over at the individual level instructions.

1. On the **Jobs** page, select **Failover** from the overflow menu for the group. Group actions will only be available when all servers in the group can safely perform that action. If you have only one server in a group, you will only have group actions.
2. You will see the failover options you specified when you created the job. If you want to edit the options, click **Edit failover settings**.



- **Source Server**—Select a server from the list and then configure the options on the rest of the page for that selected server. Repeat this process for each server in the list.
- **Authoritative failover**—This option will only be displayed if you are failing over a backup domain controller that is running in functional level 2008 R2 or older and your failover group does not contain a primary domain controller. In this specific case, you can enable

this option so that the backup domain controller will be authoritative after failover. However, you should use this option with caution. Enable it when performing a live or recovery point failover and you need to rebuild your domain in the cloud. You can also enable it when performing a test failover to an isolated network. However, disable this option when your primary domain controller is still online or you will corrupt your source domain. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.

- **Replica VM**—Specify how you want the replica virtual machine to be created in the cloud during failover.
 - **Name**—Specify the virtual machine display name. This is the replica virtual machine that will be created in the cloud. By default, this is the name of the original source server with the suffix Replica.
 - **Container Name**—Specify the name of the container to create in the cloud. If the name already exists, Carbonite Recover will append a unique number to the name.
 - **Size**—Select the size of the replica virtual machine. You can select **Specify** and identify the amount of **Memory** and the number of **Cores/Sockets** for the replica virtual machine. You may also have predefined sizes set by Carbonite. If you have predefined sizes but are uncertain what the specifications are for the size, contact Carbonite
- **Failover Type**—Specify the type of failover you want to perform.
 - **Perform failover using live data**—Select this option to initiate a full, live failover using the current data on the target appliance.
 - **Perform failover using data from a recovery point**—If you have taken recovery points, you have the option of failing over data from a recovery point. When you select this option, the list of available recovery points will appear. Select the recovery point you want to failover to. The replica data on the target appliance will be reverted to that point in time and then failover will be initiated. The **Status** and **Description** help you understand what recovery points are available.
- **Failover Network**—These are the networking values to be used on the replica virtual machine during live and recovery point failover.
 - **Network**—Select the network that you want the adapter to use on the replica virtual machine. Keep in mind the following when selecting a network.
 - If you have not been assigned two networks or you do not understand your available networks, check with Carbonite.
 - If you select a cloud network that is an isolated network, the servers that are failed over can communicate with each other but may not be able to communicate outside the isolated network. You have a few choices for working with this type of environment.
 - Include all required servers in the job group so all servers failover together.
 - Exclude a server from the job group, but add the required services from that server to another server that is in the job group.

- Use an SSL VPN client to provide a route to the isolated network. This would be a point-to-site VPN so only the servers running the SSL VPN client will have connectivity.
- **Subnet**—The subnet should default to the correct setting based on the network you selected. If needed, specify a different subnet to assign to the adapters. Make sure the subnet is valid for the selected **Network**.
- **Gateway**—The gateway should default to the correct setting based on the network you selected. If needed, specify a different gateway to assign to the adapters. Make sure the gateway is valid for the selected **Network**.
- **DNS Addresses**—The DNS addresses should default to the correct settings based on the network you selected. If needed, specify different DNS addresses to apply to the adapter. Make sure the DNS addresses are valid for the selected **Network** and your replica virtual machines will be able to reach the DNS server after failover. Click **+** to add another row to the table or **-** to remove an existing row from the table. Specify them in the order you want them used.
- **Adapter Type**—For each adapter on the source, select the network adapter type you want to use on the replica. The types available in the list will depend on the operating system you have selected, and if you have appropriate utilities, such as VMware Tools, installed on your source server.
- **IP Mode**—For each adapter on the source, select how you want the IP address to be assigned on the replica. Select **Pool** if you want the replica virtual machine to be assigned an IP address from a pool of addresses. Select **Manual** and then specify an IP address, if you want to assign a specific IP address to the adapter. The specified address must be valid for the selected **Network** and it cannot be allocated already. Click **+** to add another row to the table or **-** to remove an existing row from the table.

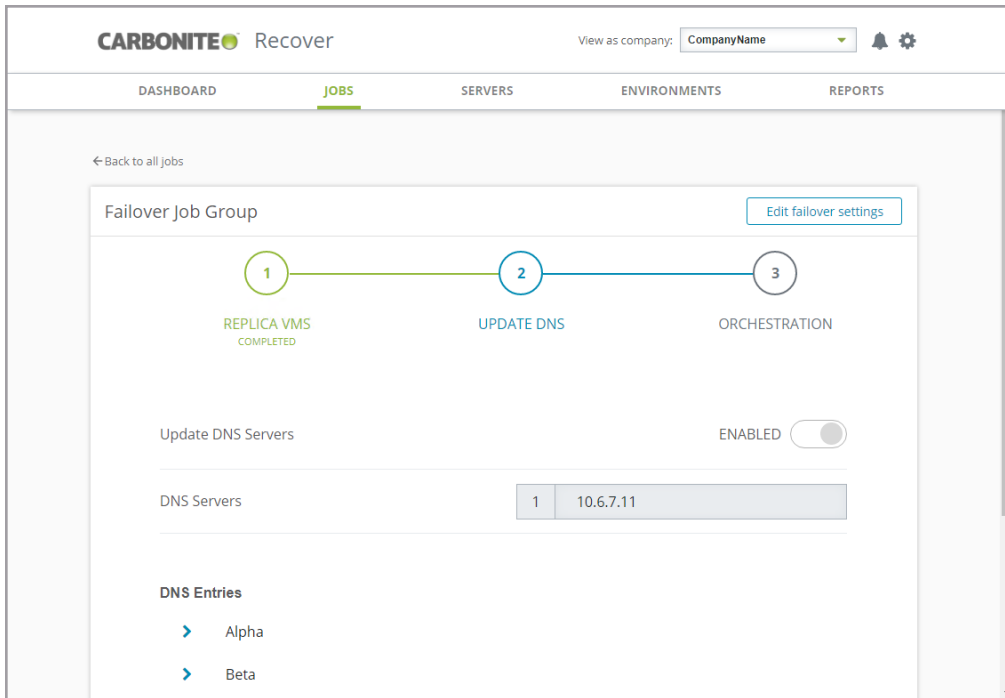


You should be protecting a source DNS server in order to be able to access other protected source servers after failover. Carbonite will provide you with an IP address to specify for your replica DNS server in the cloud. If you do not know the IP address to use, check with Carbonite.

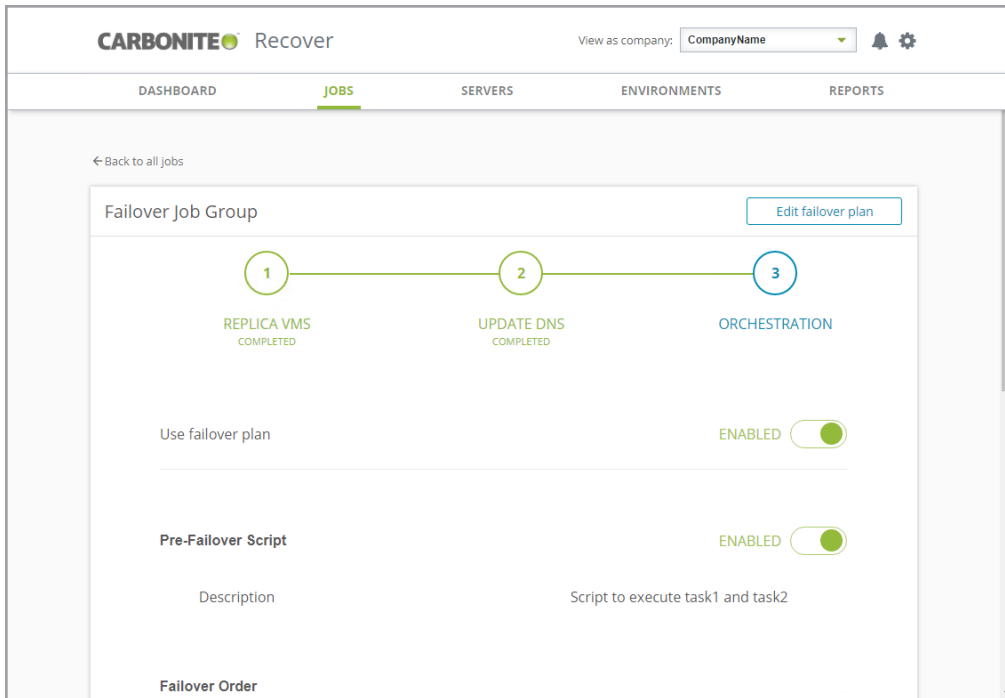
3. Click **Next** to continue.
4. Review your DNS settings. If you need to modify them, click **Edit failover settings**.



If you are only failing over Linux servers, DNS updates are not applicable. Step 2 will be for Orchestration as described in the next step of these instructions.



- **Update DNS Servers**—Enable this option if you want Carbonite Recover to update your DNS servers during failover. The toggle circle will be on the right and green when enabled and on the left and gray when disabled. This option allows you to update DNS automatically at failover time or you can trigger the updates manually. If you disable this option, you will not be able to update DNS automatically or manually.
 - **DNS Servers**—The list of DNS servers is populated from your original protection job configuration. If you did not configure DNS updates, you will see the DNS servers associated with the source servers you are protecting and the failover networks in the cloud that you have selected for each NIC on those source servers. If you do not want to update one of these DNS servers, remove it from the list by clicking the minus icon. If you want to add a DNS server that is not in the list, click the plus icon and enter the IP address.
 - **DNS Entries**— If the replica IP address after failover was configured for pool, then choose **Auto** to automatically use the assigned pool address after failover. If the replica IP address after failover was configured manually for a specific address, select the address you want DNS to use after failover. In either case, you can also set a source IP address to **Discard**.
5. Click **Next** to continue.
 6. Review your orchestration settings. If you need to modify them, click **Edit failover settings**.



- **Use failover plan**—Enable this option to start the servers in the order specified. You have the option of enabling or disabling scripts as desired. If you disable this option, the servers will all start at the same time and the scripts will be automatically disabled. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.
- **Edit plan**—This link allows you to edit any of your existing plan settings. You can change the script settings or server order.
 - **Pre-Failover Script**—Before failover starts, you can have your own script launched on a particular server. This script must be a PowerShell or Linux bash script. If you are using the failover plan, you can disable or enable scripts so they do or do not run. If you are not using the failover plan, scripts will automatically be disabled. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.
 - **Script name**—Browse (by default the local machine) and select the script that you want to run before the failover process starts. Once you select a script, the rest of the script fields will be displayed. The selected script will be encrypted and uploaded to be executed on the specified server at the designated time.
 - **Arguments**—Identify any arguments that you want passed into the script. If you are using a PowerShell script, the argument list must be a space separated list and follow standard PowerShell usage rules for hyphens, string quoting, and using special characters. For example, your PowerShell arguments might be `-parameter1 'value1' -parameter2 'value2'`. If you are using a bash script, the argument list must be a space separated list and follow standard bash usage rules for string quoting and using special

characters. For example, your bash arguments might be "value1 'value2' value3".

- **Description**—You must add a unique description to the script. The description is used to identify the script.
- **Run script on**—Select the server where you want the script to run. Keep in mind, workers cannot execute bash scripts. They can only execute PowerShell scripts. Also, if you select a source server and that server is down, the pre-failover script will not be able to be run.
- **If script fails, continue with failover**—If a script does not complete within ten minutes, the script will be considered a failure. Additionally, if there are any failures while the script is executing, the script will be considered a failure. The failover process can continue even if the script execution fails. If you disable this option, a script failure will stop the failover process. You will have to fix the script failure and restart the failover process. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.



Click **Delete script** to remove a script you have already specified.

- **Failover Order**— You can set the failover order of the servers. You will not see this section if your job group contains only one server.
 - **Use startup order for failover and fallback**—You will only see this option if you are not using scripts. Enable this option to allow for server ordering. The toggle circle will be on the right and green when enabled and on the left and gray when disabled. If you are using scripts, you will not see this option because server ordering is automatically enabled.

Drag and drop the servers in the group to the order you want them failed over. If you are using scripts, the pre-failover script will be associated with the first server in the list and the post-failover script will be associated with the last server in the list. (Associated meaning the script is executed when it is that server's turn in the server order, not that the script will run on that server.) Servers in the list will not power on until the replica virtual machine before it in the startup order has completed all of its failover process operations.



If you are protecting domain joined servers, keep in mind the following.

- You must include one or more domain controllers in your protection group so that the domain servers can resolve DNS and authenticate.
- The primary domain controller must be authoritative.
- The primary, authoritative domain controller should be a DNS server.
- You must configure failover and fallback server ordering.



- The primary, authoritative domain controller must be the first server in the server ordering.
 - All remaining domain controllers must be after the primary, authoritative domain controller and before other domain servers.
 - Make sure you add the replica networking to Windows Sites and Services.
 - If you want to remotely install the replication agent on the servers, you must use domain credentials when adding the server to Carbonite Recover. If you do not use domain credentials, you must install the replication agent on the servers manually.
-

- **Post-Failover Script**—After the failover process is completed (when the last server in the startup order is online), you can have your own script launched on a particular server. This script must be a PowerShell or Linux bash script. If you are using the failover plan, you can disable or enable scripts so they do or do not run. If you are not using the failover plan, scripts will automatically be disabled. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.
 - **Script Name**—Browse (by default the local machine) and select the script that you want to run after the failover process completes. Once you select a script, the rest of the script fields will be displayed. The selected script will be encrypted and uploaded to be executed on the specified server at the designated time.
 - **Arguments**—Identify any arguments that you want passed into the script. If you are using a PowerShell script, the argument list must be a space separated list and follow standard PowerShell usage rules for hyphens, string quoting, and using special characters. For example, your PowerShell arguments might be `-parameter1 'value1' -parameter2 'value2'`. If you are using a bash script, the argument list must be a space separated list and follow standard bash usage rules for string quoting and using special characters. For example, your bash arguments might be `"value1 'value2' value3"`.
 - **Description**—You must add a unique description to the script. The description is used to identify the script.
 - **Run Script On**—Select the server where you want the script to run. Keep in mind, workers cannot execute bash scripts. They can only execute PowerShell scripts.
-



Click **Delete script** if you need to remove a script you have already specified.

7. Click **Verify** to continue.

8. Carbonite Recover validates settings for each source server. The **Verification Checklist** page displays the validation items. Expand a source server name to see the validation items associated with that server.

Errors are designated by a white X inside a red circle. Warnings are designated by a white exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle.

Depending on the warning or error, you may see a button allowing you to **Fix** or **Fix All**. This will allow Carbonite Recover to correct the problem for you. For those warnings or errors that Carbonite Recover cannot correct automatically or any fixes that could not be successfully completed, you will need to manually correct the problem. You can revalidate the servers by clicking **Recheck**.

You can also search for specific items by using the filter.

You can continue with warnings, however, you must correct any errors before you can continue.

9. Once your configuration has passed verification with no errors, click **Finish** to start the failover.



Keep in mind, the live failover process will attempt to shut down the source server you are failing over.

If you are failing over to an isolated network and your target environment worker does not have access to that network, Carbonite Recover will not be able to verify the readiness of the replica virtual machine or complete other post-failover tasks.

Failing over at the individual level

Use this process if you are failing over an individual server from a group of multiple servers. If you are failing over a group of multiple servers or if you are failing over a group that has only one server, use the failing over at the group level instructions.

1. On the **Jobs** page, select **Failover** from the overflow menu for an individual server.
2. You will see the failover options you specified when you created the job. If you want to edit the options, click **Edit failover settings**.



If you are using scripts, the pre-failover script will be associated with the first server in the list and the post-failover script will be associated with the last server in the list. (Associated meaning the script is executed when it is that server's turn in the server order, not that the script will run on that server.) If you are failing over either the first or last server by itself, you will see a warning on the **Failover Job** page. It is a notification that scripts will not be run since you are only failing over that single server and not the group.

- **Replica VM**—Specify how you want the replica virtual machine to be created in the cloud during failover.
 - **Name**—Specify the virtual machine display name. This is the replica virtual machine that will be created in the cloud. By default, this is the name of the original source server with the suffix Replica.

- **Container Name**—Specify the name of the container to create in the cloud. If the name already exists, Carbonite Recover will append a unique number to the name.
- **Size**—Select the size of the replica virtual machine. You can select **Specify** and identify the amount of **Memory** and the number of **Cores/sockets** for the replica virtual machine. You may also have predefined sizes set by Carbonite. If you have predefined sizes but are uncertain what the specifications are for the size, contact Carbonite
- **Authoritative failover**—This option will only be displayed if you are failing over a backup domain controller that is running in functional level 2008 R2 or older and your failover group does not contain a primary domain controller. In this specific case, you can enable this option so that the backup domain controller will be authoritative after failover. However, you should use this option with caution. Enable it when performing a live or recovery point failover and you need to rebuild your domain in the cloud. You can also enable it when performing a test failover to an isolated network. However, disable this option when your primary domain controller is still online or you will corrupt your source domain. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.
- **Failover Type**—Specify the type of failover you want to perform.
 - **Perform failover using live data**—Select this option to initiate a full, live failover using the current data on the target appliance.
 - **Perform failover using data from a recovery point**—If you have taken recovery points, you have the option of failing over data from a recovery point. When you select this option, the list of available recovery points will appear. Select the recovery point you want to failover to. The replica data on the target appliance will be reverted to that point in time and then failover will be initiated. The **Status** and **Description** help you understand what recovery points are available.
- **Failover Network**—These are the networking values to be used on the replica virtual machine during live and recovery point failover.
 - **Network**—Select the network that you want the adapter to use on the replica virtual machine. Keep in mind the following when selecting a network.
 - If you have not been assigned two networks or you do not understand your available networks, check with Carbonite.
 - If you select a cloud network that is an isolated network, the servers that are failed over can communicate with each other but may not be able to communicate outside the isolated network. You have a few choices for working with this type of environment.
 - Include all required servers in the job group so all servers failover together.
 - Exclude a server from the job group, but add the required services from that server to another server that is in the job group.
 - Use an SSL VPN client to provide a route to the isolated network. This would be a point-to-site VPN so only the servers running the SSL VPN client will have connectivity.

- **Subnet**—The subnet should default to the correct setting based on the network you selected. If needed, specify a different subnet to assign to the adapters. Make sure the subnet is valid for the selected **Network**.
- **Gateway**—The gateway should default to the correct setting based on the network you selected. If needed, specify a different gateway to assign to the adapters. Make sure the gateway is valid for the selected **Network**.
- **DNS Addresses**—The DNS addresses should default to the correct settings based on the network you selected. If needed, specify different DNS addresses to apply to the adapter. Make sure the DNS addresses are valid for the selected **Network** and your replica virtual machines will be able to reach the DNS server after failover. Click **+** to add another row to the table or **-** to remove an existing row from the table. Specify them in the order you want them used.
- **Adapter Type**—For each adapter on the source, select the network adapter type you want to use on the replica. The types available in the list will depend on the operating system you have selected, and if you have appropriate utilities, such as VMware Tools, installed on your source server.
- **IP Mode**—For each adapter on the source, select how you want the IP address to be assigned on the replica. Select **Pool** if you want the replica virtual machine to be assigned an IP address from a pool of addresses. Select **Manual** and then specify an IP address, if you want to assign a specific IP address to the adapter. The specified address must be valid for the selected **Network** and it cannot be allocated already. Click **+** to add another row to the table or **-** to remove an existing row from the table.

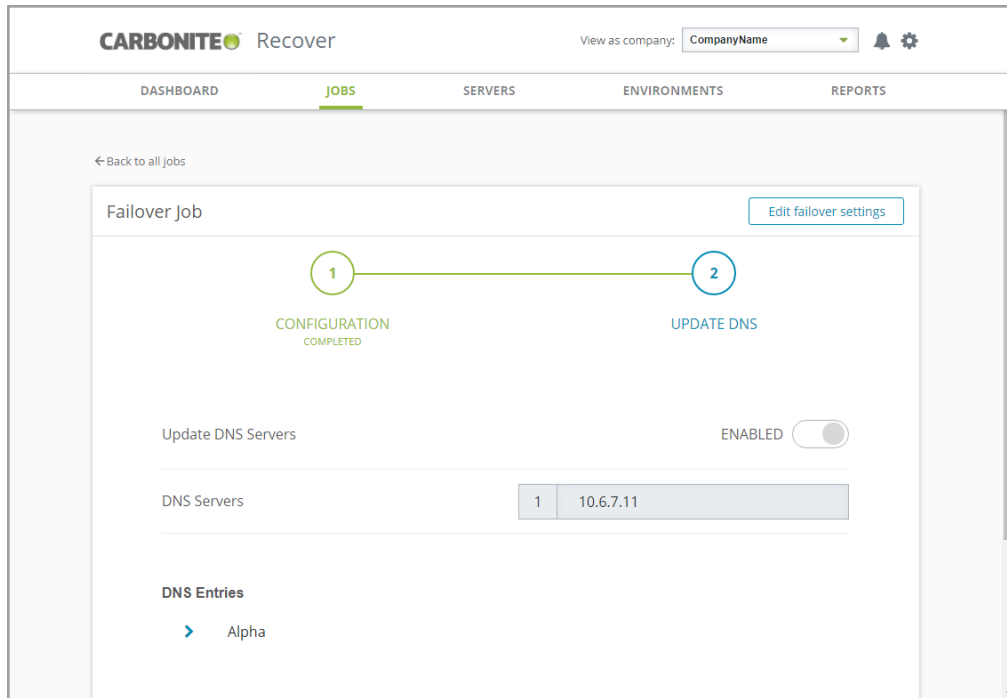


You should be protecting a source DNS server in order to be able to access other protected source servers after failover. Carbonite will provide you with an IP address to specify for your replica DNS server in the cloud. If you do not know the IP address to use, check with Carbonite.

3. Click **Next** to continue.
4. Review your DNS settings. If you need to modify them, click **Edit failover settings**.



If you are only failing over a Linux server, DNS updates are not applicable. Step 2 will not be displayed.



- **Update DNS Servers**—Enable this option if you want Carbonite Recover to update your DNS servers during failover. The toggle circle will be on the right and green when enabled and on the left and gray when disabled. This option allows you to update DNS automatically at failover time or you can trigger the updates manually. If you disable this option, you will not be able to update DNS automatically or manually.
 - **DNS Servers**—The list of DNS servers is populated from your original protection job configuration. If you did not configure DNS updates, you will see the DNS servers associated with the source servers you are protecting and the failover networks in the cloud that you have selected for each NIC on those source servers. If you do not want to update one of these DNS servers, remove it from the list by clicking the minus icon. If you want to add a DNS server that is not in the list, click the plus icon and enter the IP address.
 - **DNS Entries**— If the replica IP address after failover was configured for pool, then choose **Auto** to automatically use the assigned pool address after failover. If the replica IP address after failover was configured manually for a specific address, select the address you want DNS to use after failover. In either case, you can also set a source IP address to **Discard**.
5. Click **Verify** to continue.
 6. Carbonite Recover validates settings for each source server. The **Verification Checklist** page displays the validation items. Expand a source server name to see the validation items associated with that server.

Errors are designated by a white X inside a red circle. Warnings are designated by a white exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle.

Depending on the warning or error, you may see a button allowing you to **Fix** or **Fix All**. This will allow Carbonite Recover to correct the problem for you. For those warnings or errors that Carbonite Recover cannot correct automatically or any fixes that could not be successfully completed, you will need to manually correct the problem. You can revalidate the servers by clicking **Recheck**.

You can also search for specific items by using the filter.

You can continue with warnings, however, you must correct any errors before you can continue.

7. Once your configuration has passed verification with no errors, click **Finish** to start the failover.



Keep in mind, the live failover process will attempt to shut down the source server you are failing over.

If you are failing over to an isolated network and your target environment worker does not have access to that network, Carbonite Recover will not be able to verify the readiness of the replica virtual machine or complete other post-failover tasks.

Restoring

After you have failed over to the cloud, you can restore from the replica virtual machine back to your original source or to another server. The restore process workflow is different if you are restoring at the group level or at an individual level. The workflow also varies if you are failing over only Linux servers.

- **Group level**—Use the group level instructions if you are restoring a group of multiple servers or if you are restoring a group that has only one server.
- **Individual level**—Use the individual instructions if you are restoring an individual server from a group of multiple servers.

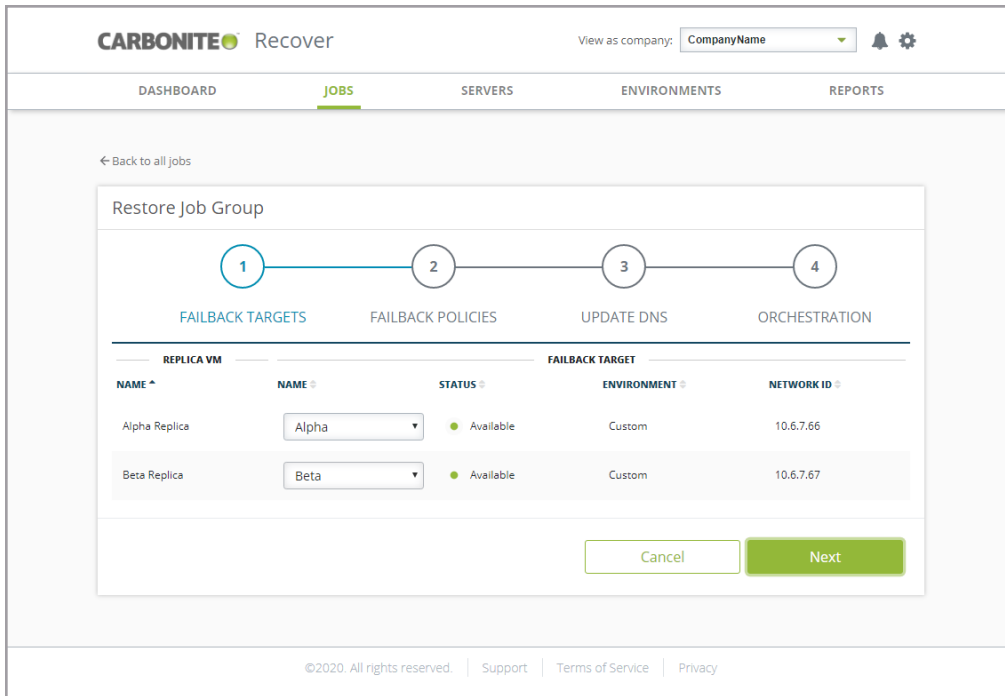
Restoring at the group level

Use this process if you are restoring a group of multiple servers or if you are restoring a group that has only one server. If you are restoring an individual server from a group of multiple servers, use the restoring at the individual level instructions.

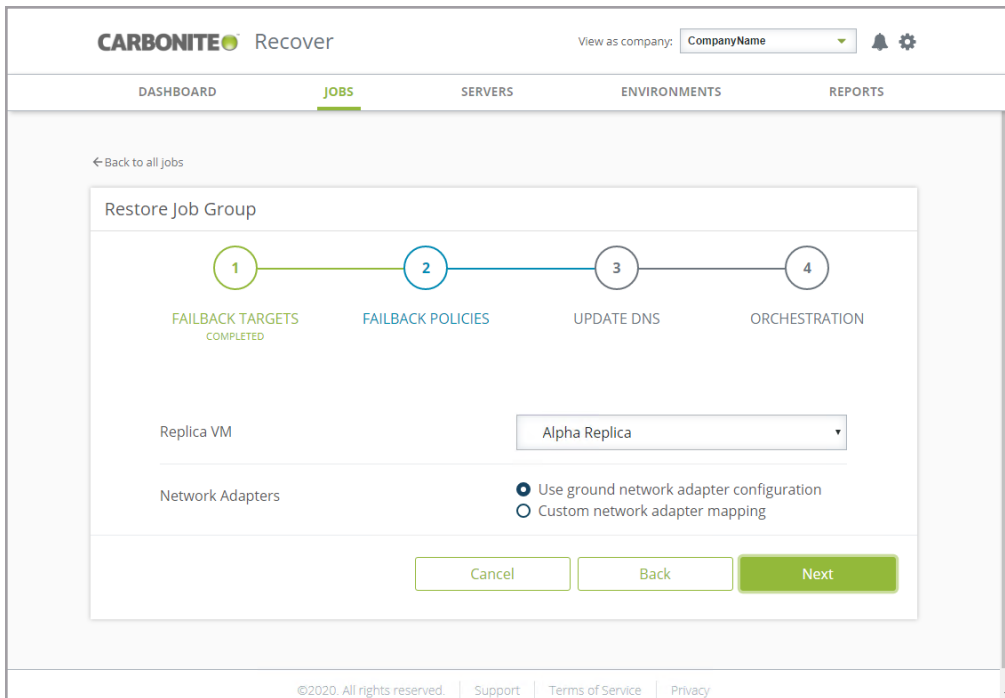
1. Your first step will depend on how you failed over your original sources to the cloud.
 - **Replica virtual machine uses same IP address as original source**—If you chose to use the same IP address on the replica virtual machine in the cloud as was used on any original source, you must bring the original source up offline, assign it a new IP address, and then bring it online. If you are restoring to a different server, make sure it has a unique IP address when it is brought online.
 - **Replica virtual machine uses different IP address than original source**—If you chose to use a different IP address on the replica virtual machine in the cloud than was used on any original source, no additional steps are required. You can bring the original source online as is or use a different server with a unique IP address.
2. On the **Jobs** page, select **Restore** from the overflow menu for the group. Group actions will only be available when all servers in the group can safely perform that action. If you have only one server in a group, you will only have group actions.
3. For each replica virtual machine you are restoring, select the failback target you want to restore to. The list of available servers will only contain those servers that are inserted in your servers list that have the replication agent installed on them and are the same operating system as your replica virtual machine which is now standing in for your original source. Your selected servers must be online before you can continue.



A domain controller cannot be failed back to an existing domain controller, including the original source domain controller. You must select an alternate server to use as the failback target for a domain controller.



4. Click **Next** to continue.
5. For each replica virtual machine you are restoring, specify how you want to handle the network adapters. Specify the network adapters for each replica virtual machine by selecting the server tabs.



- **Replica VM**—Select a replica virtual machine from the list and then configure **Network Adapters**. Repeat this process for each replica virtual machine in the list.

- **Network Adapters**—Select how you want to configure the network adapters on each failback target.
 - **Use ground network adapter configuration**—This option will leave the configuration of the network adapters on the failback target as is and use that configuration after failback.
 - **Custom network adapter mapping**—This option allows you to apply the configuration of the network adapters on the replica virtual machine to the failback target. Map the network adapters from your replica virtual machine to the adapters on the failback target server. You can also choose to **Ignore** the network adapters from your replica virtual machine. Ignoring the network adapter will not use it on the failback target. Any network adapters on the failback target server that are not mapped to a replica virtual machine adapter will be left as is.
6. Click **Next** to continue.
 7. Review your DNS settings. If you need to modify them, click **Edit DNS settings**.

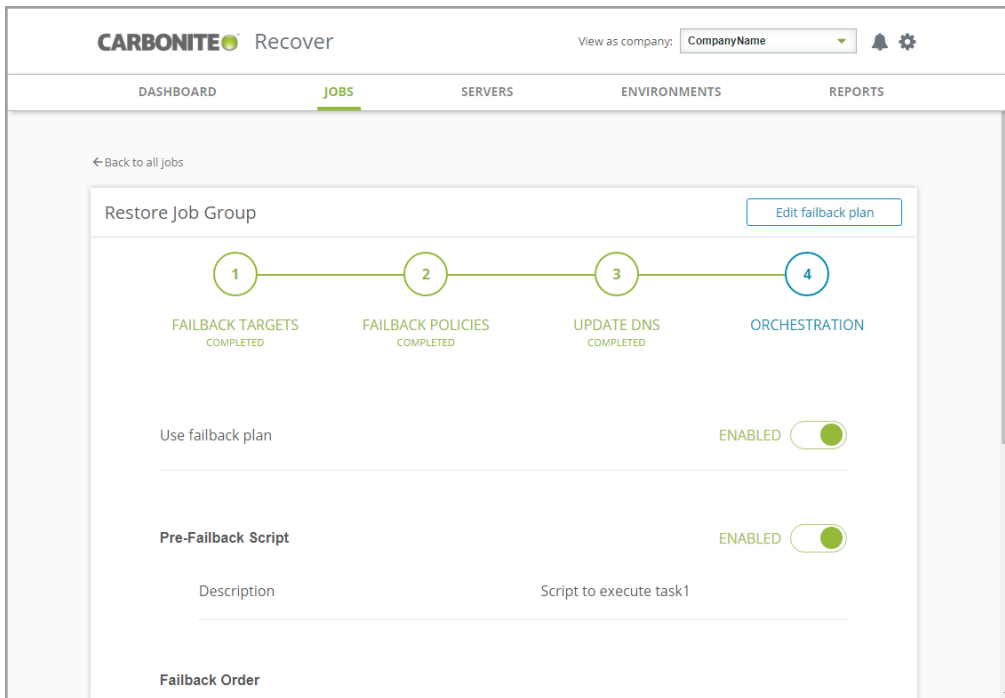


If you are only restoring Linux servers, DNS updates are not applicable. Step 3 will be for Orchestration as described in the next step of these instructions.

- **Update DNS Servers**—Enable this option if you want Carbonite Recover to update your DNS servers during failover. The toggle circle will be on the right and green when enabled and on the left and gray when disabled. This option allows you to update DNS automatically at failover time or you can trigger the updates manually. If you disable this option, you will not be able to update DNS automatically or manually.

- **DNS Servers**—The list of DNS servers is populated from the DNS servers associated with what you failed over and any additional DNS servers associated with the failback target you are using. If you do not want to update one of these DNS servers, remove it from the list by clicking the minus icon. If you want to add a DNS server that is not in the list, click the plus icon and enter the IP address.
- **DNS Entries**—For each IP address on your Windows replica virtual machine in the cloud, specify the address you want DNS to use after failback. You can also set a replica IP address to **Discard**.

8. Click **Next** to continue.
9. Review and if needed, modify your restore plan.



- **Use failback plan**—Enable this option to start the servers in the order specified. You have the option of enabling or disabling scripts as desired. If you disable this option, the servers will all start at the same time and the scripts will be automatically disabled. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.
- **Edit plan**—This link allows you to edit any of your existing plan settings. You can change the script settings or server order.
 - **Use scripts**—If you are using the failback plan, you can disable or enable scripts so they do or do not run. If you are not using the failback plan, scripts will automatically be disabled. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.
 - **Pre-Failback Script**—Before failback starts, you can have your own script launched on a particular server. This script must be a PowerShell or Linux bash script. If you are using the failback plan, you can disable or enable scripts so they do or do not run. If you are not using the failback plan, scripts will automatically be

disabled. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.

- **Script Name**—Browse (by default the local machine) and select the script that you want to run before the failback process starts. Once you select a script, the rest of the script fields will be displayed. The selected script will be encrypted and uploaded to be executed on the specified server at the designated time.
- **Arguments**—Identify any arguments that you want passed into the script. If you are using a PowerShell script, the argument list must be a space separated list and follow standard PowerShell usage rules for hyphens, string quoting, and using special characters. For example, your PowerShell arguments might be `-parameter1 'value1' -parameter2 'value2'`. If you are using a bash script, the argument list must be a space separated list and follow standard bash usage rules for string quoting and using special characters. For example, your bash arguments might be `"value1 'value2' value3"`.
- **Description**—You must add a unique description to the script. The description is used to identify the script.
- **Run Script On**—Select the server where you want the script to run. Keep in mind, workers cannot execute bash scripts. They can only execute PowerShell scripts.
- **If script fails, continue with restore**—If a script does not complete within ten minutes, the script will be considered a failure. Additionally, if there are any failures while the script is executing, the script will be considered a failure. The failback process can continue even if the script execution fails. If you disable this option, a script failure will stop the failback process. You will have to fix the script failure and restart the failback process. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.



Click **Delete Script** if you need to remove a script you have already specified.

- **Failback Order**— Your servers are currently listed in the failover order you configured. If desired, drag and drop the servers in the group to the order you want them failed back. Servers in the list will not start the failback process until the server before it in the startup order has completed all of its failback process operations.



If you are protecting domain joined servers, keep in mind the following.

- You must include one or more domain controllers in your protection group so that the domain servers can resolve DNS and authenticate.



- The primary domain controller must be authoritative.
- The primary, authoritative domain controller should be a DNS server.
- You must configure failover and failback server ordering.
- The primary, authoritative domain controller must be the first server in the server ordering.
- All remaining domain controllers must be after the primary, authoritative domain controller and before other domain servers.
- Make sure you add the replica networking to Windows Sites and Services.
- If you want to remotely install the replication agent on the servers, you must use domain credentials when adding the server to Carbonite Recover. If you do not use domain credentials, you must install the replication agent on the servers manually.

-
- **Post-Failback Script**—After the failback process is completed (when the last server in the startup order is online), you can have your own script launched on a particular server. This script must be a PowerShell or Linux bash script. If you are using the failback plan, you can disable or enable scripts so they do or do not run. If you are not using the failback plan, scripts will automatically be disabled. The toggle circle will be on the right and green when enabled and on the left and gray when disabled.
 - **Script Name**—Browse (by default the local machine) and select the script that you want to run after the failback process completes. Once you select a script, the rest of the script fields will be displayed. The selected script will be encrypted and uploaded to be executed on the specified server at the designated time.
 - **Arguments**—Identify any arguments that you want passed into the script. If you are using a PowerShell script, the argument list must be a space separated list and follow standard PowerShell usage rules for hyphens, string quoting, and using special characters. For example, your PowerShell arguments might be `-parameter1 'value1' -parameter2 'value2'`. If you are using a bash script, the argument list must be a space separated list and follow standard bash usage rules for string quoting and using special characters. For example, your bash arguments might be `"value1 'value2' value3"`.
 - **Description**—You must add a unique description to the script. The description is used to identify the script.
 - **Run Script On**—Select the server where you want the script to run. Keep in mind, workers cannot execute bash scripts. They can only execute PowerShell scripts.



Click **Delete script** if you need to remove a script you have already specified.

10. Click **Verify** to continue.
11. Carbonite Recover validates settings for each replica virtual machine. The **Verification Checklist** page displays the validation items. Expand a replica virtual machine name to see the validation items associated with that server.

Errors are designated by a white X inside a red circle. Warnings are designated by a white exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle.

Depending on the warning or error, you may see a button allowing you to **Fix** or **Fix All**. This will allow Carbonite Recover to correct the problem for you. For those warnings or errors that Carbonite Recover cannot correct automatically or any fixes that could not be successfully completed, you will need to manually correct the problem. You can revalidate the servers by clicking **Recheck**.

You can also search for specific items by using the filter.

You can continue with warnings, however, you must correct any errors before you can continue.

12. Once your configuration has passed verification with no errors, click **Finish** to start the restoration.

Restoring at the individual level

Use this process if you are restoring an individual server from a group of multiple servers. If you are restoring a group of multiple servers or if you are restoring a group that has only one server, use the restoring at the group level instructions.

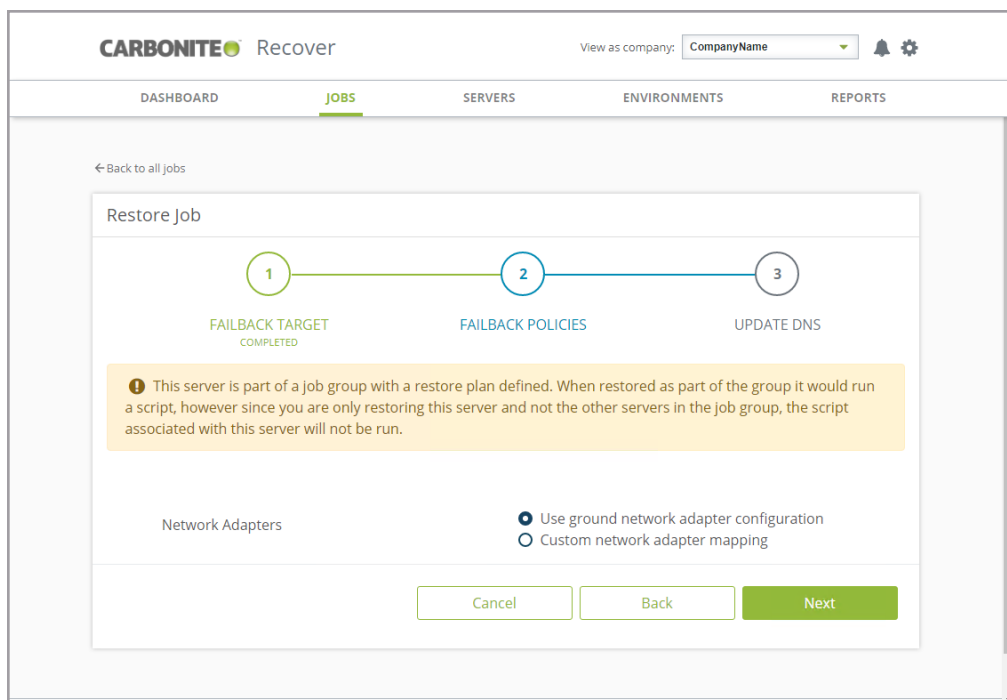
1. Your first step will depend on how you failed over your original sources to the cloud.
 - **Replica virtual machine uses same IP address as original source**—If you chose to use the same IP address on the replica virtual machine in the cloud as was used on any original source, you must bring the original source up offline, assign it a new IP address, and then bring it online. If you are restoring to a different server, make sure it has a unique IP address when it is brought online.
 - **Replica virtual machine uses different IP address than original source**—If you chose to use a different IP address on the replica virtual machine in the cloud than was used on any original source, no additional steps are required. You can bring the original source online as is or use a different server with a unique IP address.
2. On the **Jobs** page, select **Restore** from the overflow menu for an individual server.
3. For the replica virtual machine you are restoring, select the failback target you want to restore to. The list of available servers will only contain those servers that are inserted in your servers list and are the same operating system as your replica virtual machine which is now standing in for your original source. Your selected server must be online before you can continue.

©2020. All rights reserved. | Support | Terms of Service | Privacy

4. Click **Next** to continue.
5. Specify how you want to handle the network adapters.



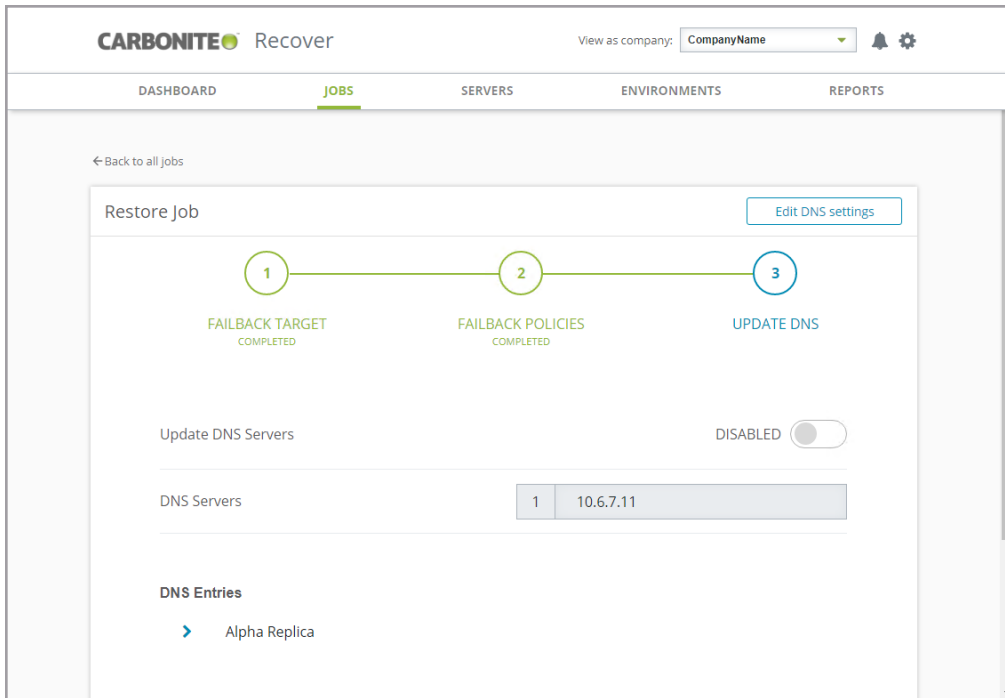
If you are using scripts, the pre-failback script will be associated with the first server in the list and the post-failback script will be associated with the last server in the list. (Associated meaning the script is executed when it is that server's turn in the server order, not that the script will run on that server.) If you are failing back either the first or last server by itself, you will see a warning on the **Restore Job** page. It is a notification that scripts will not be run since you are only failing back that single server and not the group.



- **Use ground network adapter configuration**—This option will leave the configuration of the network adapters on the failback target as is and use that configuration after failback.
 - **Custom network adapter mapping**—This option allows you to apply the configuration of the network adapters on the replica virtual machine to the failback target. Map the network adapters from your replica virtual machine to the adapters on the failback target server. You can also choose to **ignore** the network adapters from your replica virtual machine. Ignoring the network adapter will not use it on the failback target. Any network adapters on the failback target server that are not mapped to a replica virtual machine adapter will be left as is.
6. Click **Next** to continue.
 7. Review your DNS settings. If you need to modify them, click **Edit DNS settings**.



If you are only restoring a Linux server, DNS updates are not applicable. Step 3 will not be displayed.



- **Update DNS Servers**—Enable this option if you want Carbonite Recover to update your DNS servers during failover. The toggle circle will be on the right and green when enabled and on the left and gray when disabled. This option allows you to update DNS automatically at failover time or you can trigger the updates manually. If you disable this option, you will not be able to update DNS automatically or manually.
 - **DNS Servers**—The list of DNS servers is populated from the DNS servers associated with what you failed over and any additional DNS servers associated with the failback target you are using. If you do not want to update one of these DNS servers, remove it from the list by clicking the minus icon. If you want to add a DNS server that is not in the list, click the plus icon and enter the IP address.
 - **DNS Entries**—For each IP address on your Windows replica virtual machine in the cloud, specify the address you want DNS to use after failback. You can also set a replica IP address to **Discard**.
8. Click **Verify** to continue.
 9. Carbonite Recover validates settings for each replica virtual machine. The **Verification Checklist** page displays the validation items. Expand a replica virtual machine name to see the validation items associated with that server.

Errors are designated by a white X inside a red circle. Warnings are designated by a white exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle.

Depending on the warning or error, you may see a button allowing you to **Fix** or **Fix All**. This will allow Carbonite Recover to correct the problem for you. For those warnings or errors that Carbonite Recover cannot correct automatically or any fixes that could not be successfully completed, you will need to manually correct the problem. You can revalidate the servers by clicking **Recheck**.

You can also search for specific items by using the filter.

You can continue with warnings, however, you must correct any errors before you can continue.

10. Once your configuration has passed verification with no errors, click **Finish** to start the restoration.

Failing back

After the restoration is complete, you need to complete failback. This finalizes the identity transfer from the replica virtual machine to the failback target.

On the **Jobs** page, select **Failback** from the overflow menu for an individual server or a group. Group actions will only be available when all servers in the group can safely perform that action. If you have only one server in a group, you will only have group actions.

When the failback is complete, the failback target will be a replica of your replica virtual machine in the cloud, including any changes that were made to that replica virtual machine while it was running in the cloud.

You can reprotect the failback target again by selecting **Reprotect** from the overflow menu. You can reuse the hard disks that were created during the last job. You can also create new disks, if desired. In either case, the replica virtual machine created in the cloud from the last job will be deleted. Keep in mind, any recovery points taken prior to the failover will not be available when you reprotect.

Viewing job details

On the **Jobs** page, in the overflow menu on the right of a group level table row, select **View details**. On the **Viewing** page, you will find group statistics as well as individual server details and statistics.



You will see details while a job is in the process of being deleted, but once it is deleted, the details will no longer be available.

The screenshot shows the Carbonite Recover web interface. At the top, there's a navigation bar with 'CARBONITE Recover' on the left and 'View as company: CompanyName' on the right. Below the navigation bar are tabs for 'DASHBOARD', 'JOBS', 'SERVERS', 'ENVIRONMENTS', and 'REPORTS'. The 'JOBS' tab is active. The main content area shows 'Viewing Alpha and Beta' with a list of groups: 'Alpha' and 'Beta'. A 'Group Statistics' table is displayed on the left, and a 'Job Details' panel is open on the right. The 'Job Details' panel shows 'Source Server: Alpha', 'Target Appliance: WindowsAppliance', and 'Status: Protecting'. Below the 'Job Details' panel are two sub-panels: 'Operations' and 'Statistics'. The 'Operations' panel has an 'Add Job Group' button. The 'Statistics' panel shows a table with columns for 'Activity' and 'Connected Since'.

Group Statistics	
Disk Queue	-
Initial Mirror Complete	true
SSH Direct Connect	false
Mirror Remaining	-
Mirror Skipped	7.4 GB
Recovery Point Latency	0 seconds
Replication Queue	-
Data Sent	52.2 GB
Compressed Data Sent	52.2 GB

Job Details	
Source Server	Alpha
Target Appliance	WindowsAppliance
Status	Protecting

Statistics	
Activity	Protecting
Connected Since	04/16/2020 22:45:23
Disk Queue	-
Initial Mirror Complete	true

- **Group Statistics**—These statistics are cumulative for all of the jobs in the group.
 - **Disk Queue**—This is the amount of disk space being used to queue data on the source servers (when protecting) or on the replica virtual machines (when restoring).
 - **Initial Mirror Complete**—This field indicates if all of the initial copies of data have completed from your source servers to the target appliances (when protecting) or from the replica virtual machines to the failback targets (when restoring).
 - **SSH Direct Connection**—This field will always be false.
 - **Mirror Remaining**—This is the amount of data remaining to be sent from the source servers to the target appliances (when protecting) or from the replica virtual machines to the failback targets (when restoring).
 - **Mirror Skipped**—This is the amount of data that has been skipped because the data is not different on the source servers and target appliances (when protecting) or on the replica virtual machines and failback targets (when restoring).

- **Recovery Point Latency**—This is the longest length of time replication is behind on any one target appliance compared to the source server they are protecting or on any one failback target compared to the replica virtual machine they are restoring from. This is the longest 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 synchronization data. If you are synchronizing and failover (or synchronizing and failback), the data on the target appliance (or the failback target) will be at least as far behind as the replication point latency. It could potentially be further behind depending on the circumstances of the synchronization. If synchronization is idle and you failover (or failback), the data will only be as far behind as the replication point latency time.
- **Replication Queue**—This is the amount of disk space being used to queue replication data on the source servers (when protecting) or replica virtual machines (when restoring).
- **Data Sent**—This is the total amount of data sent from the source servers to the target appliances (when protecting) or from the replica virtual machines to the failback targets (when restoring).
- **Compressed Data Sent**—This is the total amount of compressed data sent from the sources servers to the target appliances (when protecting) or from the replica virtual machines to the failback targets (when restoring). If compression is disabled, this statistic will be the same as bytes sent.
- **Job Details**—Click a server name above **Group Statistics** to see server details, operations, statistics, and recovery points for that individual server.
 - **Source Server**—During the protecting and failover states, the source of the job is your source server. During the restoring and failback states, the source of the job is replica virtual machine.
 - **Target Appliance**—During the protecting and failover states, the target of the job is your target appliance. During the restoring and failback states, the target of the job is the failback target.
 - **Status**—The status indicates, by color and description, the health of the individual job.
 - **Green**—A green circle indicates a good status.
 - **Yellow**—A yellow circle indicates a pending or warning status. Generally, Carbonite Recover is working, waiting on a pending process, or attempting to resolve the warning state.
 - **Red**—A red circle indicates an error status. You will need to investigate and resolve the error.
 - **Black**—A black circle indicates the status is unknown.



Additional statuses are coming directly from the replication agent and can provide further information when a job is in an error state.

- **Operations**—This section shows the operations being performed for the individual server. You can expand sub-sections to see the specific tasks within an operation. The operations list is in chronological order.

- **Statistics**—This section shows the statistics for the individual server.
 - **Activity**—This field indicates any activity information coming from the replication agent. When this field is blank, it indicates there is no activity from the replication agent. For example, when you are failing over, there is no more data being replicated between the source and target appliance.
 - **Connected Since**—When you first create your job, this is the date and time when the disks on the target appliance are first attached. If you have started a stopped job, this is the date and time when the job was restarted.
 - **Disk Queue**—This is the amount of disk space being used to queue data on the source servers (when protecting) or on the replica virtual machines (when restoring).
 - **Initial Mirror Complete**—This field indicates if all of the initial copies of data have completed from your source servers to the target appliances (when protecting) or from the replica virtual machines to the failback targets (when restoring).
 - **SSH Direct Connection**—This field will always be false.
 - **Mirror Started**—This is the time the most recent synchronization started.
 - **Mirror Finished**—This is the time the most recent synchronization ended. If this field is blank, synchronization is currently in progress.
 - **Mirror Remaining**—This is the amount of data remaining to be sent from the source servers to the target appliances (when protecting) or from the replica virtual machines to the failback targets (when restoring).
 - **Mirror Skipped**—This is the amount of data that has been skipped because the data is not different on the source servers and target appliances (when protecting) or on the replica virtual machines and failback targets (when restoring).
 - **Mirror State**—This field indicates the status of synchronization. (Replication of data changes are on-going and continuous.)
 - **Recovery Point Latency**—This is the longest length of time replication is behind on any one target appliance compared to the source server they are protecting or on any one failback target compared to the replica virtual machine they are restoring from. This is the longest 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 synchronization data. If you are synchronizing and failover (or synchronizing and failback), the data on the target appliance (or the failback target) will be at least as far behind as the replication point latency. It could potentially be further behind depending on the circumstances of the synchronization. If synchronization is idle and you failover (or failback), the data will only be as far behind as the replication point latency time.
 - **Replication Queue**—This is the amount of disk space being used to queue replication data on the source servers (when protecting) or replica virtual machines (when restoring).
 - **Data Sent**—This is the total amount of data sent from the source servers to the target appliances (when protecting) or from the replica virtual machines to the failback targets (when restoring).
 - **Compressed Data Sent**—This is the total amount of compressed data sent from the sources servers to the target appliances (when protecting) or from the replica virtual

machines to the failback targets (when restoring). If compression is disabled, this statistic will be the same as bytes sent.

- **Recovery points**—If there are any recovery points for the server, they will be listed. Newest recovery points will be at the top of the list. You can use recover points to failover to an earlier point in time. To help you understand what recovery points are available, the **Description** indicates if the recovery point was scheduled, manually taken (taken on demand), or skipped (the initial mirror was not yet complete). If you no longer need a recovery point, select the overflow menu at the right of a table row and select **Delete** to remove it from the job.

Chapter 8 Configuring email notification

By default, you will receive email messages for the notifications generated by Carbonite Recover. Use the following instructions to change your email notification settings.

1. Click on the gear icon in the upper right corner of the Carbonite Recover web page, and select your user name.
2. On the **Preferences** tab, modify your email settings as needed.
 - **Subject Prefix**—By default, the subject line of email alerts sent to your account email address will be prefaced with Carbonite Recover Notification and the company name. This prefix allows you to recognize and filter emails specific to Carbonite Recover and companies. You can change or remove the first part of prefix as desired (not the company name). The remainder of the subject line will contain the notification content, truncated if necessary. The email body will contain the full notification content.
 - **Notifications**—Select the type and level of notifications that you want to receive as email messages. If you do not select any type or level, you will not receive notifications as email messages. You will still get notifications in the Carbonite Recover web interface whether email notifications are enabled or disabled.
3. Click **Save**.



Time references in your email notifications will be in UTC time.

Chapter 9 Viewing reports and company usage

You can view failover and usage reports to help you understand your protections, failovers, and storage consumption. The reports are available per company. If you are assigned to multiple companies, select the company that you want to view the report for by selecting the company name in the **View as company** drop-down list next to the bell notification icon. Select a company from the list or filter the list by typing in text and then selecting a company from the filter.

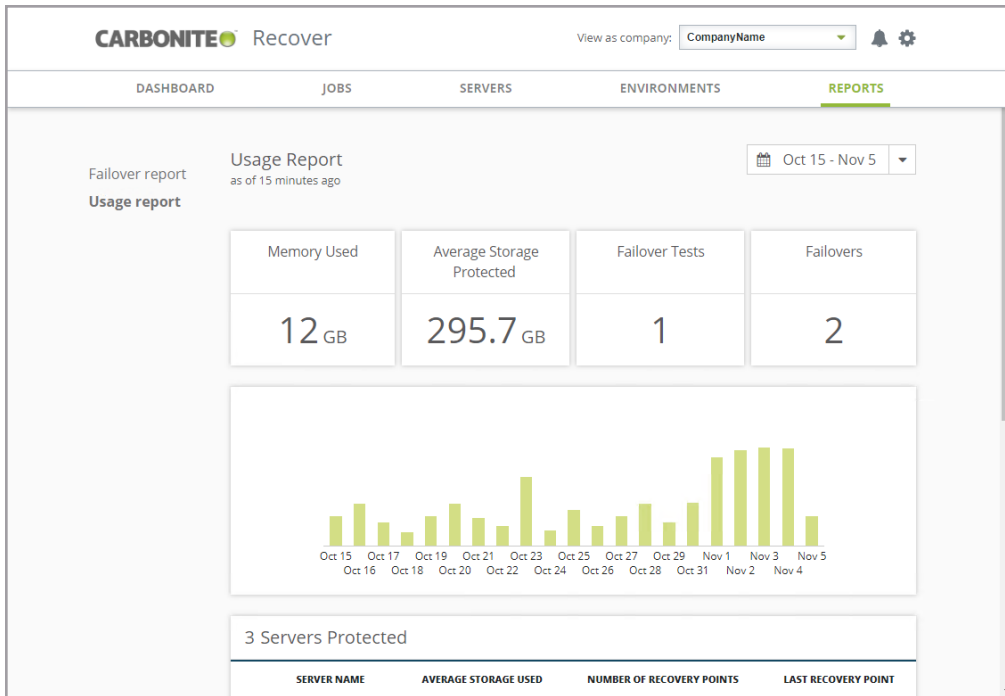
- **Failover report**—On the **Reports** page on the **Failover report** tab, you will see a list of completed failovers, both test and live, for the selected company.

FAILOVER DATE	TYPE	JOB GROUP	TIME TO FAILOVER	OUTCOME	SUMMARY
Nov 5, 2019 3:58 PM	Test	Delta	4 Minutes	Successful	
Nov 4, 2019 8:20 AM	Failover	Alpha and Beta	13 Minutes	Successful	Self-entered description

- **Sort**—Sort the table by clicking a column heading. When the arrow is pointing up, the table is sorted by that column in ascending order. When the arrow is pointing down, the table is sorted by that column in descending order.
- **Filter**—Text entered in a filter box and selected from a filter drop-down list will narrow the list displayed to only those rows that contain the search text and selected item.
- **High level information**—The main table rows contain high level information for each failover.
 - **Failover Date**—This is the date and time when failover started.
 - **Type**—The type of failover is live or test.
 - **Job Group**—This is the name of the job group.
 - **Time to Failover**—This is the amount of time it took for the failover process to complete.
 - **Outcome**—By default, the result of the failover as determined by Carbonite Recover is displayed. Click this hyperlink to set an **Actual Outcome** which is the outcome that you determine. Your outcome may be the same or different than the

Carbonite Recover outcome. For example, failover may have completed successfully, but if you did not update DNS, you may want to consider the result a failure.

- **Summary**—These are your own notes describing the failover. You can enter notes independent of setting an **Actual Outcome**. The text is limited to 256 characters with the first 100 characters showing on the **Reports** page.
- **Detailed information**—Expand a main table row to see detailed information for that failover.
 - **Job**—This is the name of the job within the job group.
 - **Time to Failover (RTO)**—This is the amount of time it took for the failover process to complete. This is often referred to as the recovery time objective.
 - **Time in Cloud**—This is the amount of time the replica virtual machine was active in the cloud. It is the difference between the start and end times.
 - **Failover Start**—This is the date and time when failover started.
 - **Failover End**—This is the date and time when failover ended. This is when undo finished for a test failover or when failback started for a live failover.
 - **Recovery Point (RPO)**—This is the recovery point used for the failover. This is often referred to as the recovery point objective. For live and test failover, this is when the failover process started. For recovery point failover, this is the time the recovery point was taken.
 - **Errors**—If a number is displayed, that is the number of errors that occurred during failover. You may have had errors even if failover was overall successful. For example, DNS could not be updated. Hover over the error number to see the tasks that failed. If nothing is displayed in this column of the table, then the job completed failover with no errors.
- **Usage report**—On the **Reports** page on the **Usage report** tab, you will see usage information for the selected company.



- **Time period**—Select a time period from the drop-down list in the upper right corner of the **Usage report** tab. Billing periods are from the 15th of one month to the 14th of the next month. You can select a specific billing period or select **Custom Range** to choose specific dates.
- **High level statistics**—At the top of the page, you will find high level statistics for your company.
 - **Memory Used**—This is the amount of memory used for the last day of the selected billing period.
 - **Average Storage**—This is the average amount of storage used for the selected billing period.
 - **Failover Tests**—This is the number of test failovers that were completed during the selected billing period.
 - **Failovers**—This is the number of live or recovery point failovers that were completed during the selected billing period.
- **Daily table**—The table below the high level statistics shows each day in the selected billing period that consumed storage in the cloud. Hover over a bar in the table to see the amount of storage consumed on that day. If there was no storage consumed in the cloud for a day, that day will not show in the table.
- **Servers Protected**—This table breaks down the storage data by server. You can also see the number of recovery points for the server and the date of the last recovery point. If you expand the server row by clicking on the right arrow to the left of the server name, you can see storage data by date.
- **Failovers**—This table shows the completed failovers (live, recovery point, or test). You can see when the replica virtual machine was powered on and off along with memory usage for the replica virtual machine. Failed failovers (live, recovery point, or test) will not appear on the usage report and are not billed.

Chapter 10 Administration

If you were assigned the Administrator role for your company, you will be able to perform additional tasks that the User role does not have access to.

- *Managing users* on page 95—Company administrators can add, view, edit, or delete users within their company.
- *Managing workers* on page 101—If Carbonite did not install a worker for you, or you need to install one in your source environment, a company administrator can install, view, and delete workers for their company.

Managing users

If you were assigned the Administrator role for your user account, an additional tab in the user interface is available. The **Users** tab allows you to create users and assign them to your company and your child companies.



You cannot modify your own account.

The **Users** page provides high-level information and controls for your users for the currently selected company.

<input type="checkbox"/>	NAME ^	ROLE	EMAIL	STATUS	COMPANY
	Filter by name	All	Filter by email	All	All
	FirstName LastName	Company Admin	name@domain.com	Enabled	CompanyName
<input type="checkbox"/>	User2 LastName2	User	name2@domain.com	Enabled	CompanyName
<input type="checkbox"/>	User3 LastName3	User	name3@domain.com	Enabled	CompanyName

The following controls are available on the **Users** page.

- **Add user**—Click this button to add users to a company. See *Adding a user* on page 97 for details.
- **Table overflow menu**—When at least one row in the table is selected, you will have an overflow menu at the top of the **Users** page. Select **Delete** to delete the selected users. Since jobs are associated with a company, not the user who created it, jobs will continue to function after a user is deleted.
- **Select All** and **Clear All**—Click the checkbox in the column heading to toggle between selecting all items on that page of the table or clearing all selections on that page of the table. This option will not apply to items on a page that are hidden by a search filter.
- **Sort**—Sort the table by clicking a column heading. When the arrow is pointing up, the table is sorted by that column in ascending order. When the arrow is pointing down, the table is sorted by that column in descending order.

- **Filter**—Text entered in a filter box and selected from a filter drop-down list will narrow the list displayed to only those rows that contain the search text and selected item.
- **Status**—This column has different colors and status information to indicate the health of your users.
 - **Green**—A green circle indicates a confirmed, active user.
 - **Yellow**—A yellow circle indicates a user that has not yet confirmed the invitation. The user is inactive.
 - **Red**—A red circle indicates a disabled user or a user with an error. You will need to enable the user or investigate and resolve the error.
- **Table row overflow menu**—In the overflow menu on the right of a table row, you can select the following actions.
 - **View details**—Select this option to view or edit the user details. See *Viewing and editing user details* on page 99 for details.
 - **Invite**—Select this option to send another invitation to a user who has not yet confirmed an invitation.
 - **Disable**—Select this option to disable this user's access. The company association will still be defined, but the user will not have access to Carbonite Recover. Since jobs are associated with a company, not the user who created it, jobs will continue to function after a user is disabled.
 - **Enable**—Select this option to enable access to Carbonite Recover.
 - **Delete**—Select this option to delete the user. Since jobs are associated with a company, not the user who created it, jobs will continue to function after a user is deleted.
- **Table paging**—At the bottom of the table you will see the row numbers you are currently viewing and the total number of table rows. Paging buttons allow you to move between pages of the table. The single arrow buttons move forward or backward one page. The double arrow buttons move to the first or last page.

Adding a user

When you add a new user, that user will receive a registration email. The user's status remains inactive until they confirm the registration email, when their status changes to enabled.

1. Select the **Users** tab.
2. Click **Add user** and enter the user information.

The screenshot shows the 'Add User' form in the Carbonite Recover interface. The form is titled 'Add User' and is located within the 'USERS' tab. It contains three main input fields: 'Company', 'Email', and 'Role'. The 'Company' field is a dropdown menu with a search filter 'Filter by company name' and a selected option 'CompanyName'. The 'Email' field is a text input with the value 'UserName@domain.com' and a required field asterisk. The 'Role' field has two radio button options: 'User' (selected) and 'Administrator'. At the bottom of the form are 'Cancel' and 'Save' buttons. The footer of the application shows '©2020. All rights reserved. | Support | Terms of Service | Privacy'.

- **Company**—Select the company the user should have access to.
- **Email**—Enter the email address for the user.
- **Role**—Specify the user's role.
 - **User**—This role type can manage environments, servers, and jobs in the assigned company and all of the child companies.
 - **Administrator**—This role type can create new users in the assigned company and all of the child companies. This user can also manage environments, servers, and jobs in the assigned company and in all of the child companies.

The following table summarizes the tasks each type of user can perform. Where access is granted, the access is at the user's assigned company and all child companies.

Task	User Type	
	Administrator	User
Manage users	✓	

Manage environments	✓	✓
Manage servers	✓	✓
Manage jobs	✓	✓
See company reports and usage	✓	✓
Enable email notifications	✓	✓

3. Click **Save**. The invitation email is sent to the user's email address. The user will not be active until the email invitation is confirmed.

Viewing and editing user details

You can edit a user's account information (name or role), but not the user's email address. The email address is tied directly to the account. When you edit a user, no additional registration email is sent. The modified changes take effect immediately.

1. On the **Users** page, find the table row of the user you want to view or edit. In the overflow menu for that table row, click **View details**.
2. On the **User details** page, view or modify the user information.

- **Company**—Select the company the user should have access to.
- **Name**—Enter the first and last name of the user.
- **Email**—This field is read-only and cannot be changed.
- **Role**—Specify the user's role.
 - **User**—This role type can manage environments, servers, and jobs in the assigned company and all of the child companies.
 - **Administrator**—This role type can create new users in the assigned company and all of the child companies. This user can also manage environments, servers, and jobs in the assigned company and in all of the child companies.

The following table summarizes the tasks each type of user can perform. Where access is granted, the access is at the user's assigned company and all child companies.

Task	User Type	
	Administrator	User

Manage users	✓	
Manage environments	✓	✓
Manage servers	✓	✓
Manage jobs	✓	✓
See company reports and usage	✓	✓
Enable email notifications	✓	✓

3. Click **Save** to save any changes to the user details.
4. On the **Delete** tab, you can delete the user. Since jobs are associated with a company, not the user who created it, jobs will continue to function after a user is deleted.

Managing workers

If Carbonite did not install a worker for you, or you need to install one in your source environment, a company administrator can install, view, and delete workers for their company. A company administrator can also register any server where you will be running PowerShell scripts that need to communicate with Carbonite Recover



An account assigned the User role for the company can view workers and script clients but cannot install, register, or delete them.

1. Go to **Settings > Manage instance** and click the **Workers and Clients** tab.
2. Select the company you want to view from the **View as company** drop-down list. This will change the displayed clients and workers to only those for that selected company.
3. To install a worker, click **Download**. See *Installing a worker* on page 103 for details.
4. To generate a registration token to be used in a PowerShell script, click **Generate Registration Token**. You must use the registration token within ten minutes of creation.
5. In the **Registered Workers** and **Registered Clients** sections, you can manage your servers that have been registered as workers or registered to run PowerShell scripts.

The screenshot shows the Carbonite Recover web interface. At the top, there's a header with the Carbonite logo and 'Recover' text. To the right, it says 'View as company: CompanyName' with a dropdown arrow, a bell icon, and a gear icon. Below the header is a navigation bar with tabs: DASHBOARD, JOBS, SERVERS, ENVIRONMENTS, USERS, and REPORTS. The main content area is titled 'Manage instance' and has a sub-section 'Workers and Clients'. Under 'Workers and Clients', there's a 'Download Worker' section. It contains a description: 'A worker is a server that receives and executes tasks on the servers and jobs in your environments. Download and install the worker on a server that has access to your environments.' Below the description, it says 'Current worker version: 2.3.0.512' and a green 'Download' button. Below the 'Download Worker' section is a 'Registered Workers' section. It features a table with the following columns: NAME, STATUS, SERVER, IP ADDRESSES, and VERSION. The table has a search bar for 'Filter by name' and dropdowns for 'All' under STATUS and IP ADDRESSES, and 'Filter by server' and 'Filter by IP addr' for the SERVER and IP ADDRESSES columns respectively. There is also a 'Filter by name' button. The table contains one row: Worker_Server_Comp..., Running, ServerName, 172.29.41.221, 2.3.0.512. At the bottom of the table, it says '1 - 1 of 1' with navigation arrows.

- **Delete**—When at least one row in the table is selected, you can select **Delete** in the overflow menu at the top right corner of a section to remove that worker or script client. It will no longer be able to communicate with the Orchestrator server after it is deleted. Once you have deleted a worker in Carbonite Recover, you can uninstall the worker software on that server.

- **Select All** and **Clear All**—Click the checkbox in the column heading to toggle between selecting all items on that page of the table or clearing all selections on that page of the table. This option will not apply to items on a page that are hidden by a search filter.
- **Sort**—Sort the table by clicking a column heading. When the arrow is pointing up, the table is sorted by that column in ascending order. When the arrow is pointing down, the table is sorted by that column in descending order.
- **Filter**—Enter text in the filter box to narrow the list to only rows that contain the search text.
- **Table row overflow menu**—In the overflow menu on the right of a table row, you can select **Delete** to remove that worker or script client. It will no longer be able to communicate with the Orchestrator server after it is deleted. Once you have deleted a worker in Carbonite Recover, you can uninstall the worker software on that server.
- **Table paging**—At the bottom of the table you will see the row numbers you are currently viewing and the total number of table rows. Paging buttons allow you to move between pages of the table. The single arrow buttons move forward or backward one page. The double arrow buttons move to the first or last page.



If you want to use a worker server again after you have deleted it from the registered workers list, you must you must uninstall and reinstall the software on the server. See *Installing a worker* on page 103 for details.

Installing a worker

If Carbonite did not install a worker for you, or you need to install one in your source environment, use the following instructions to install the worker.



Worker installations are only available to users that are assigned the Administrator role for their company.

1. From the Carbonite Recover web portal, log in using your company administrator account.
 2. Select the company for which you want to create a worker from the **View as company** drop-down list.
 3. Select **Manage instance** from the gear drop-down menu.
 4. On the **Workers and Clients** tab, click **Download**.
 5. Save the file when prompted.
-



The downloaded worker installation file is a personalized file specific to the company currently selected in the **View as company** list. Do not share this downloaded file between companies.

The worker installation file is time-sensitive. It must be used within 10 minutes of downloading, otherwise the installation will fail. If the file is more than 10 minutes old, you must download a new worker installation file.

6. Run the file.

After a few minutes, the installation will be complete. You can confirm the installation by checking for **Carbonite Recover Worker** in **Program and Features** or **Carbonite DRaaS Worker** in **Administrative Tools > Services**.



If you want to delete a worker, you must delete it from Carbonite Recover and then you can uninstall the worker software on that server. See *Managing workers* on page 101 for details on how to delete a worker.
