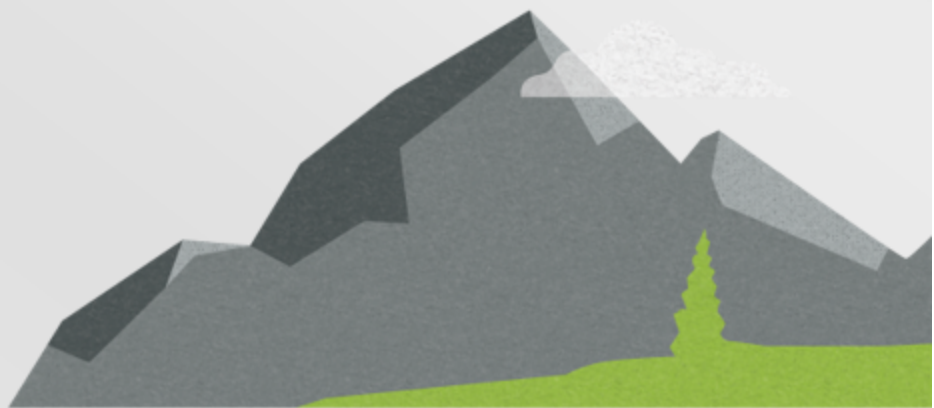# CARBONITE

Carbonite Recover

User's Guide

Notices

Carbonite Recover User's Guide Wednesday, December 20, 2017

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

Man pages are installed and available on Carbonite Availability Linux servers. These documents are bound by the same Carbonite license agreement as the software installation.

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

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

Contents

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 server in the cloud with minimal downtime. See *How Carbonite Recover works* on page 5 for a workflow of the Carbonite Recover process.

Physical   **Virtual**   **Cloud**
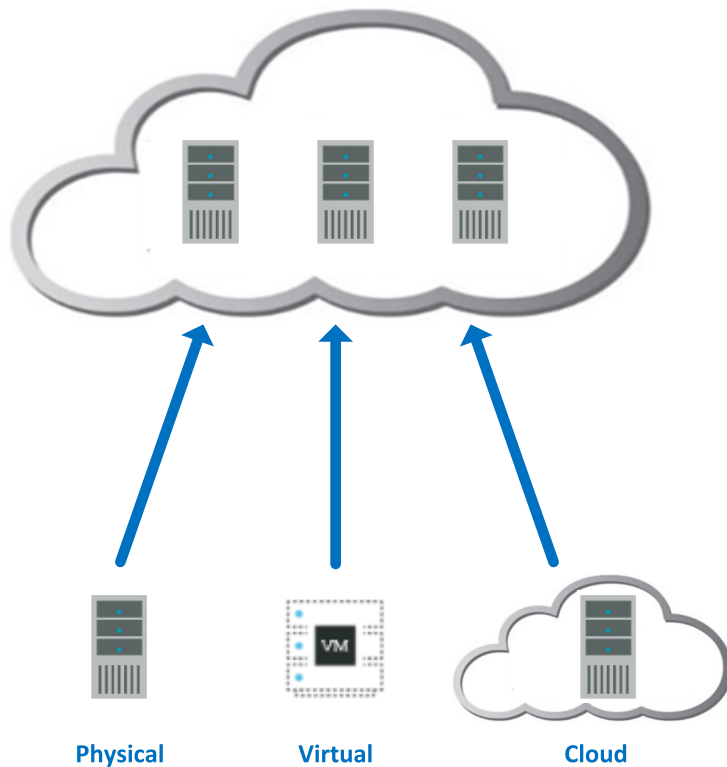
How Carbonite Recover works

Before you begin with Carbonite Recover, you have servers you want to protect to the cloud. These servers are called your source servers.
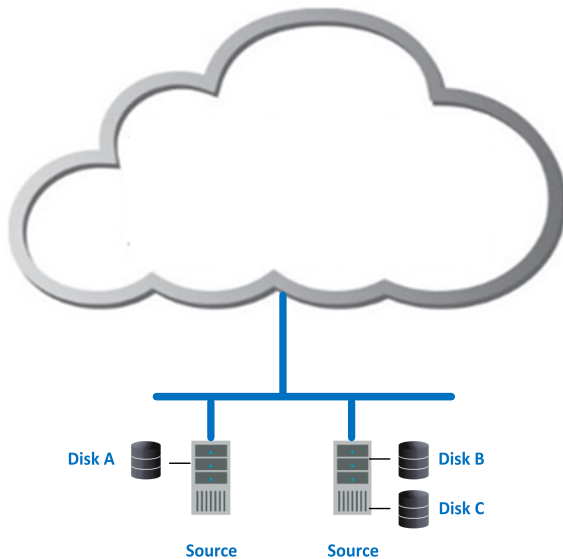
To protect your source servers, you must have at least one target appliance and at least one worker. Your service provider will create a worker for you, but you will need to create the target appliance.

- **Target appliance**—A target appliance is a virtual server in the cloud created from a template provided by Carbonite. You must have at least one target appliance, and it 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 server in the cloud.

- **Workers**—Workers receive and execute tasks by communicating with the Carbonite Recover backend infrastructure that Carbonite is running.
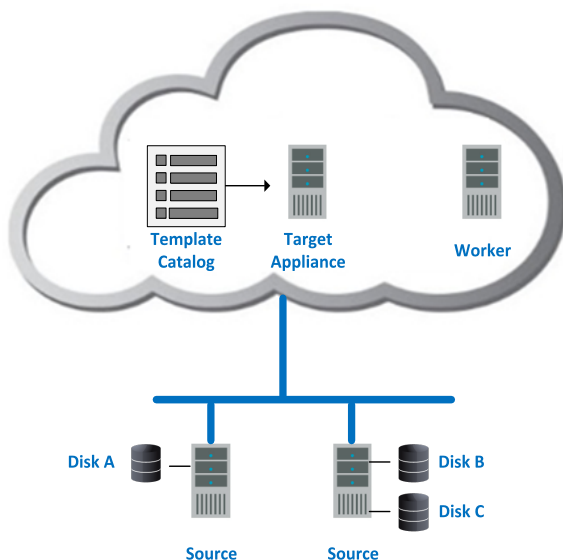
When protection begins, the target appliance maintains a replica of the data from the source servers you are protecting by using virtual hard disks attached to the target appliance.
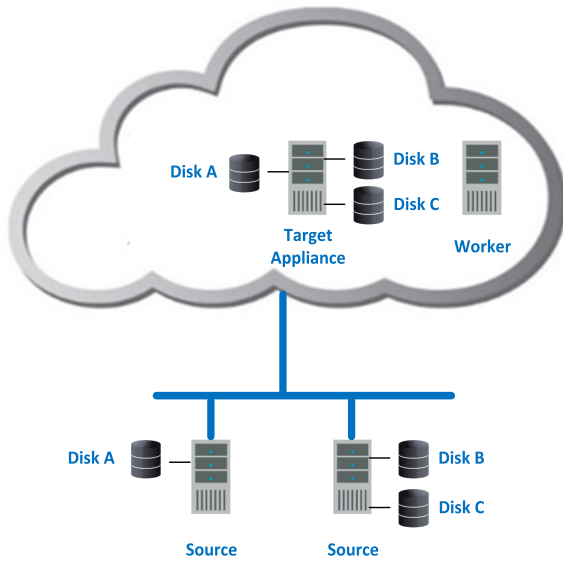
In the event a source server fails, the worker quickly creates a replica server in the cloud and detaches the hard disks from the target appliance and attaches them to the new replica server.
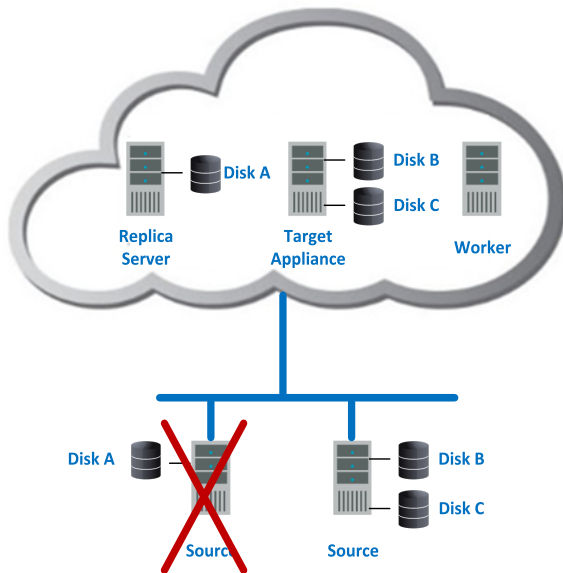
You can run on the replica server in the cloud as long as needed. When you are ready, you can restore and failback from the replica server in the cloud back to your original server or to a different server, as needed.

Chapter 2 **Requirements**

Your environment must meet the following requirements.

- **Source servers**—The source servers you are protecting must meet the following requirements.
  - **Operating system**—The servers must be one of the following operating systems running the NTFS file system.
    - 2008 R2 Service Pack 1 or later
    - 2012
    - 2012 R2

    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.

  - **System memory**—At least 1 GB of memory is required on the servers.
  - **Server name**—Unicode file system support is included, but your server name must still be in ASCII format. All servers must have a unique server name.
  - **VMware Tools**—All servers hosted on VMware must have VMware Tools installed.
  - **Windows Remote Management**—In order for Carbonite Recover to push the necessary Recover replication agent software to your source servers, you must enable Windows Remote Management (WinRM). To enable it, 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.
  - **Firewall configuration**—Each source server must have the following ports open.
    - **Windows Remote Management (HTTP-In)**—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**.
    - **Recover replication agent**—Ports 6320 and 6325 must be open for UDP and TCP for both inbound and outbound traffic.
  - **SAN policy**—Some versions and editions of Windows, like 2008 Enterprise and Datacenter and some 2012 editions, set the default SAN policy to disabled. This keeps the non-boot volumes from mounting on the replica server after failover. To avoid this issue, enable your SAN policy 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.
- **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 5.5 or later
  - **Hyper-V**—You can use any of the following Hyper-V versions.
    - 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
- **Workers**—Carbonite will create and register a worker server for you.
- **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.
- **Networking**—You must establish a VPN between the source servers and the target appliance. Also, if they are on different subnets you will need name resolution.
- **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.
- **Snapshots**—Microsoft Volume Shadow Copy service (VSS) is used for snapshot capability. Snapshots are taken and stored with the replica data on the target appliance, so make sure that you configure the target appliance large enough to maintain the source servers and snapshots.

  There are limitations imposed by VSS that impact Carbonite Recover snapshots. For example, VSS only maintains 512 snapshots. If the maximum number of snapshots exists and another one is taken, the oldest snapshot is deleted to make room for the new one.

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

  You can use VSS on your source servers for other uses outside Carbonite Recover, for example Microsoft Backup uses it. Keep in mind though that the driver for VSS is started before the driver for Carbonite Recover. Therefore, if you use snapshots on your source servers and you revert any files on the source server, Carbonite Recover will not be aware of the revert and the file change will not be replicated to the target appliance. The file change will be mirrored to the target appliance during the next mirroring process.
- **Web browser**—You will need a modern, recent web browser to access the Carbonite Recover web interface. Google Chrome and Mozilla Firefox are the preferred browsers, however Microsoft Internet Explorer version 11 or later is supported.

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 7 for details on these components. You will also need a web browser to access the Carbonite Recover interface.

Even though data is encrypted between the source server and target appliance, a VPN is required between these servers.
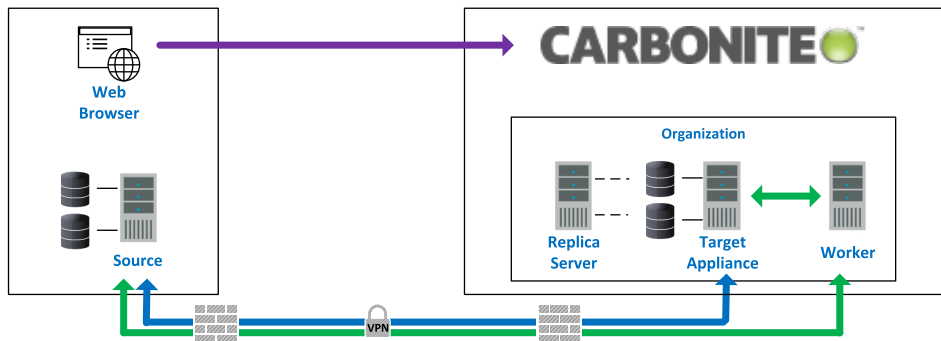
Component to Component	Communication and Port	Arrow Color
Web Browser to Carbonite Recover	HTTPS port 443	
Source Server to Worker Target Appliance to Worker	HTTPS port 6326 HTTP port 5985 and HTPPS 5986	
Source Server to Target Appliance	Recover replication agent ports 6320 and 6325	

Chapter 3 Getting started

Before you get started, make sure you have reviewed the Carbonite Recover *Requirements* on page 7. Then complete the following tasks, in order.

1. **Accept invitation**—Carbonite will send you an email invitation to access Carbonite Recover. See *Carbonite Recover interface* on page 11 for more details on accepting your invitation and an overview of the Carbonite Recover interface.

2. **Add environments** —You must create an environment for your source servers and one for the cloud. 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 15.

   Carbonite may have pre-populated the target environment for you.

3. **Add servers to your source environment**—Once you have your source environment created, you need to add servers to it 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 21.

4. **Create a target appliance in your cloud environment**—Once you have your target environment created, you need to create a target appliance in it. A target appliance is a virtual server in the cloud created from a template provided by Carbonite. You must have at least one target appliance, and it 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 server in the cloud. To create a target appliance, see *Creating a target appliance* on page 23. If you already have an existing target appliance, see *Adding servers or existing target appliances* on page 21.

5. **Create a job** —Once your source and target environments are prepared with sources and your target appliance, you can protect those sources. See *Protecting servers* on page 29 for complete details.

Chapter 4 Carbonite Recover interface

- **First time access**—Carbonite will send you an email invitation to access Carbonite Recover. When you receive the email, click **Verify Email** and you will be taken to the Carbonite Recover interface which is hosted by Carbonite. Enter your first and last name. Enter and confirm a password. Agree to the license agreement and then click **Register** to finalize your registration.

- **Bookmark**—Once you have accessed the interface, you should bookmark the page so you can easily return to the URL.

- **Dashboard**—Each time you log in to Carbonite Recover, you will see a dashboard page, which is the highest level overview. In each of the sections on the dashboard, you will see a breakdown for that particular Carbonite Recover component. You can click a tile under **Job Status** or a hyperlink next to a pie chart to jump to that page of the interface. A filter will automatically be applied to the page you jump to showing only the components that match the job status or hyperlink status you selected.
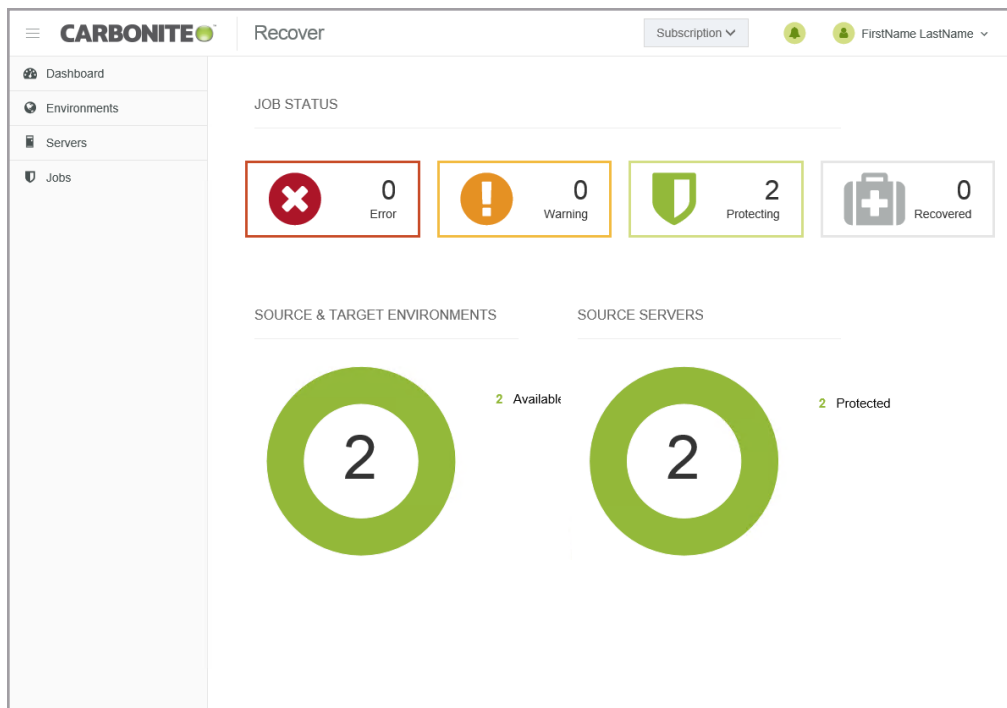
- **Subscriptions**— You can access your subscriptions from the drop-down list to the left of the bell notification icon. Your Carbonite Recover registration is associated with one or more subscriptions. Subscriptions are created by Carbonite to determine which areas of the backend infrastructure users can access. You will belong to one or more subscriptions based on subscription assignments. You will only be able to access the cloud resources assigned to the subscription you are currently using.

- **Notifications**—No matter which page you are on in the interface, you can access a list of notifications from the bell icon near the upper right corner of the page. If you have only information notifications, the circle and bell will be green. If you have one or more warning notifications, the circle and bell will be yellow. If you have one or more error notifications, the circle and bell will be

red. The yellow warning color overrides green information, and red error overrides yellow warning. By clicking the circle or bell, you can view the notifications and the approximate time they were generated. You can dismiss individual notifications or all of the notifications in the current list.

- **User profile**—Under your sign in name, in the upper right corner, is a menu of options.

    - **Subscription Usage**—This option shows your usage and billing information. See *Subscription usage* on page 57 for more details.
    - **User Preferences**—This option allows to set preferences specific to your Carbonite Recover account.

        - **Language**—Currently English is the only available language.
        - **Theme**—Select the color scheme for the Carbonite Recover interface.
        - **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.

        - **Subject Prefix**—By default, the subject line of email alerts sent to your account email address will be prefaced with Carbonite Recover Notification. This prefix allows you to recognize and filter emails specific to Carbonite Recover. You can change or remove the prefix as desired. The remainder of the subject line will contain the 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.

    - **Knowledge Base**—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.
    - **Log Out**—You can log out of the Carbonite Recover interface by clicking your account name in the upper right corner of the page and clicking **Log out**. 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.

Chapter 5 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.

- **Microsoft Hyper-V**—This environment contains a host and source servers you want to protect.
- **VMware vSphere**—This environment contains a host and source servers you want to protect.
- **Custom**—This environment contains source 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 a server hosted in Microsoft Azure.
- **VMware vCloud**—This environment contains a cloud host, the target appliance, and a worker.

On the **Environments** page, you will find high-level information and controls for your environments. You can sort the environment list by clicking on any column heading.

You can select multiple environments on this page by selecting the checkbox to the left of a table row. A checkmark indicates the environment is selected. You can also select the checkbox column heading to quickly select or deselect all environments on the current page.
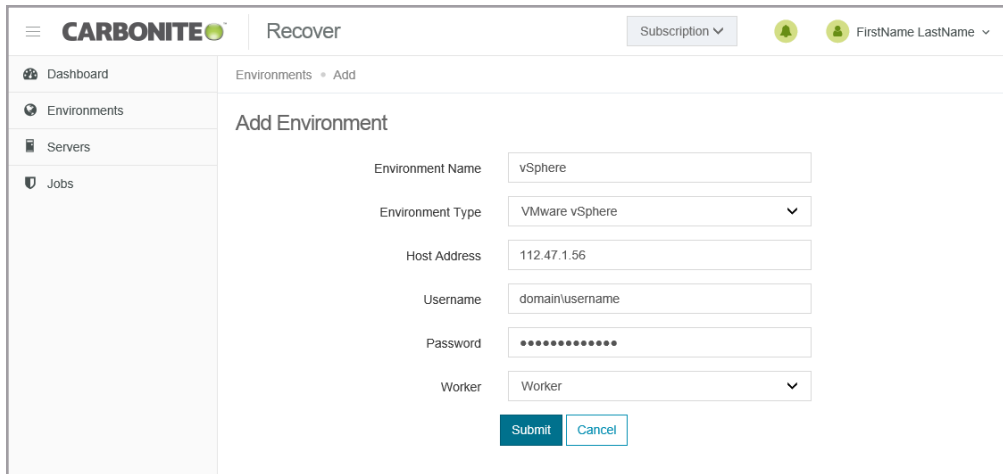

- **Toolbar and table controls**—The following controls are available on the toolbar on the **Environments** page.
    - **Add**—Click this button to add a new environment. See *Adding an environment* on page 15 for more details on this process.
    - **Delete**—Click this button to delete the selected environments. You cannot delete an environment that has established jobs.
    - **Select All**—Click this link to select all items in the table. This option will select all table rows across all pages, even if hidden by a search filter. If you want to select only the rows that are visible on the current page, select the checkbox column heading.

- **Clear All Selections**—Click this link to deselect all items in the table. This option will deselect all table rows across all pages, even if hidden by a search filter. If you want to deselect only the rows that are visible on the current page, deselect the checkbox column heading.

- **per page**—Select the number of table rows to display per page.

- **Starting typing to filter**—Text entered in the filter box will narrow the list displayed to only those rows that contain the search text.

- **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.

- **Hyperlink control**—Click the name of an environment to see the details for that environment. See *Viewing environment details* on page 16 for more information.

- **Environment status**—The **Status** column has different colors and status information to indicate the health of your 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.

- **Single environment controls**—The following menu options are available in the **Actions** menu.

  - **Add Server**—Select this option to add servers or target appliances to the environment. See *Adding servers or existing target appliances* on page 21 for more details.
  - **Delete**—Select this option to delete the environment. You cannot delete an environment that has an established job.

Adding an environment

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

- **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.
- **Username**—For Microsoft Hyper-V and VMware vSphere environments, specify a user name with access to the host. For VMware vCloud environments, 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 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, you will not see this option.
- **Worker**—Select a worker for this environment.

3. When you have identified your environment, click **Submit**.

Viewing environment details

On the **Environments** page, click on the name of an environment to see details for that specific environment.
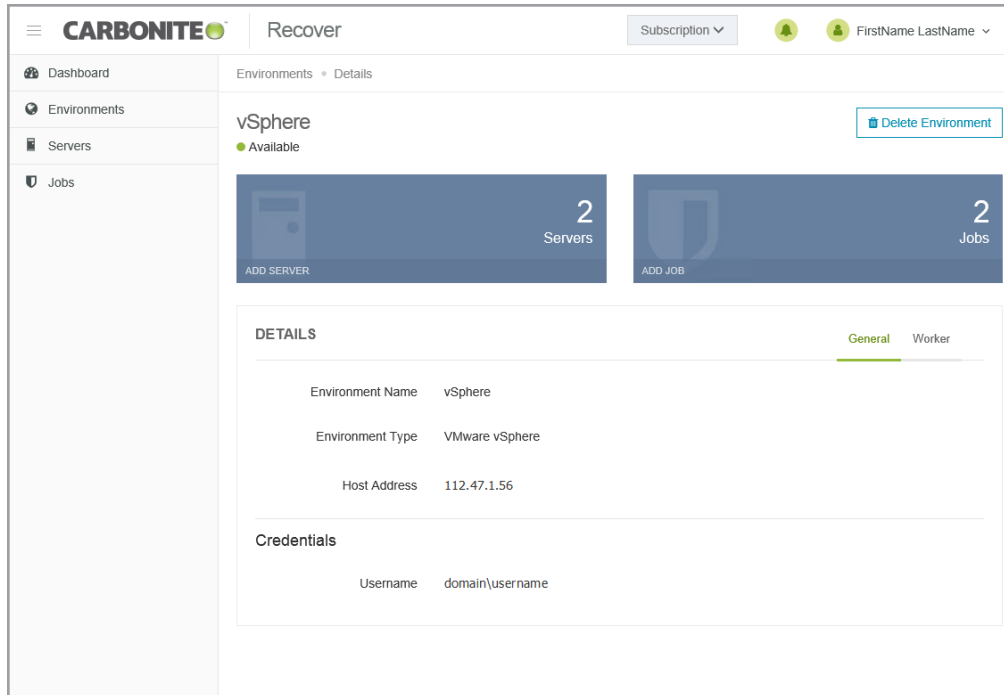
- **Toolbar controls**—There is one toolbar control available on the **Environments Details** page. Click the **Delete** button to delete the environment. You cannot delete an environment that has established jobs.

- **Environment status**—The status is displayed under the environment name and indicates, by color and description, the health of your 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.

- **Tiles**—Tiles near the top of the details page show how many of each component are associated with the environment.

  - **Servers**—This is the number of servers and target appliances in the environment. Click **Add Server** to add servers to this environment. See *Adding servers or existing target appliances* on page 21 for more details.
  - **Jobs**—This is the number of jobs in the environment. This number is not the job groups, but the individual jobs within groups. Click **Add Job** to create a new job. See *Protecting servers* on page 29 for more details.

- **Tabs**—There are two tabs below the tiles.

- **General**—This tab show the details for the environment as well as credentials. (Custom environments do not have credentials.) You can edit the environment name or credentials by hovering over an item and clicking in the field. Click **Save** to save any changes.
- **Worker**—This tab shows the worker associated with the environment. If you need to change workers, you will have to delete the environment and re-create it. Environments cannot be deleted if they have established jobs.

Chapter 6 **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.

On the **Servers** page, you will find high-level information and controls for your servers. You can sort the server list by clicking on any column heading.

You can select multiple servers on this page by selecting the checkbox to the left of a table row. A checkmark indicates the server is selected. You can also select the checkbox column heading to quickly select or deselect all servers on the current page.
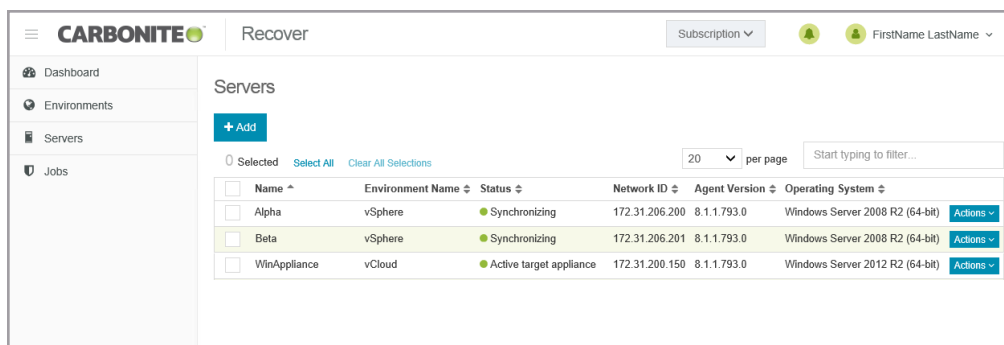
- **Toolbar and table controls**—The following controls are available on the toolbar on the **Servers** page.
  - **Add**—Click this button to add servers to an environment. See *Adding servers or existing target appliances* on page 21 for more details on adding servers or target appliances that have already been created. If you need to create a target appliance, see *Creating a target appliance* on page 23.
  - **Protect**—Click this button to protect the selected source servers. This option is only available for source servers in an unprotected state. See *Protecting servers* on page 29 for more details on this process.
  - **Refresh**—Click this button to re-inventory the selected servers and gather information from them.
  - **Change Credentials**—Click this button to update the credentials used to access the selected servers.
  - **Remove**—Click this button to remove the selected servers from your list. You cannot remove a server that has established jobs.
  - **Select All**—Click this link to select all items in the table. This option will select all table rows across all pages, even if hidden by a search filter. If you want to select only the rows that are visible on the current page, select the checkbox column heading.
  - **Clear All Selections**—Click this link to deselect all items in the table. This option will deselect all table rows across all pages, even if hidden by a search filter. If you want to deselect only the rows that are visible on the current page, deselect the checkbox column heading.

- **per page**—Select the number of table rows to display per page.
- **Starting typing to filter**—Text entered in the filter box will narrow the list displayed to only those rows that contain the search text.
- **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.
- **Server status**—The **Status** column has different colors and status information to indicate the health of your servers. 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.
    - **Black**—A black circle indicates the status is unknown.
- **Single server controls**—The following menu options are available in the **Actions** menu.
    - **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 *Protecting servers* on page 29 for more details on this process.
    - **Install Replication Agent**—Select this option to install the Recover replication agent on the server. The Recover replication agent is the engine that powers mirroring and replication from your source server to your target appliance in the cloud and from your replica server in the cloud to your failback source during restoration. This is an optional installation because the protection process will automatically install the Recover replication agent on the source during job creation, if it is not already installed.
    - **Update Credentials**—Select this option to update the credentials used to access the server.
    - **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 vCloud 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 vCloud environment.
    - **Power On**—Select this option to power on the server. This option is only available for servers running in your vCloud 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 vCloud environment.

- **Delete VM**—Select this option to delete the server from your vCloud environment. This option is only available for servers running in your vCloud environment.

Adding servers or existing target appliances

Use these instructions to add a server or a target appliance that has already been created to an environment. (If you need to create a target appliance, see *Creating a target appliance* on page 23.)

1. Create an environment, if you do not have one already. See *Adding an environment* on page 15.
2. On the **Servers** tab, click **Add**.
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, and then click **Next**.
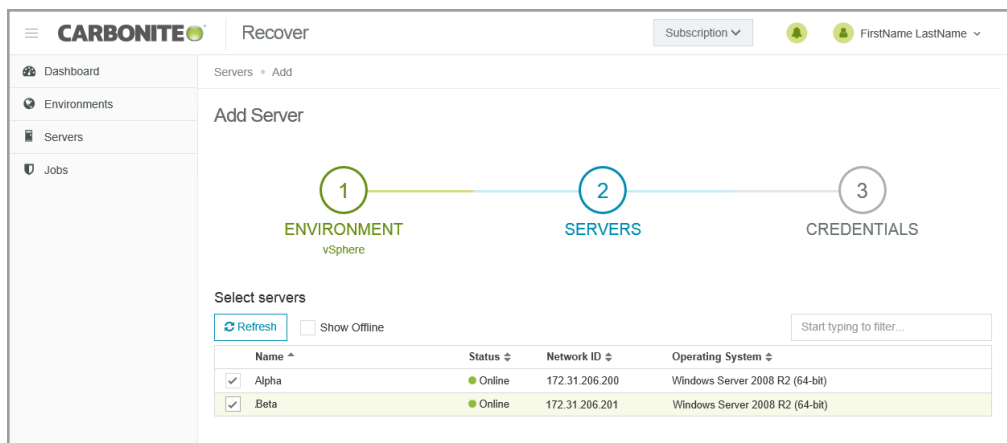
> Only online servers will be displayed by default. If you want to see all servers, select **Show Offline**. You cannot protect an offline server.
>
> Your servers must have at least one NIC attached to the server in order for the server to be discoverable.
>
> Make sure all enabled NICs are operational. An enabled but unplugged NIC will cause a server to fail to be added to your environment. Disabled NICs are not an issue.

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

2. Click **Add Server** to add another row to the table or **-** to remove an existing row from the table.
3. Click **Finish** to add the servers to the custom environment.

Creating a target appliance

A target appliance is a virtual server in the cloud created from a template provided by Carbonite. You must have at least one target appliance, and it 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 server 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 21.)

1. On the **Servers** tab, click **Add**.
2. On the **Add Server** page, confirm your vCloud environment is selected and then select **Create Appliance**.
3. Click **Next** to continue.
4. Specify the storage and platform for the target appliance.
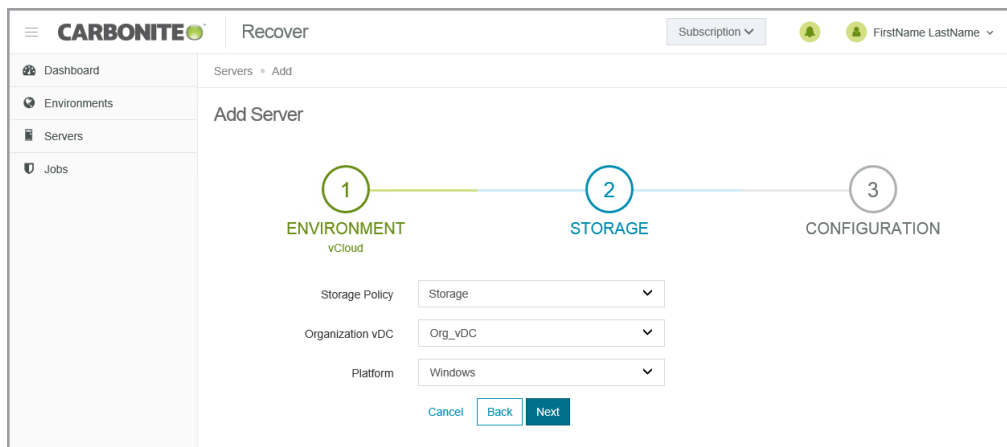
- **Storage Policy**—Select the storage policy to use from the cloud environment you selected. If you do not know your storage policies, check with Carbonite.
- **Organization vDC**—Select the organization vDC to use from the storage policy selected. If you do not know your organization vDC, check with Carbonite.
- **Platform**—Windows is the only platform currently supported.

5. Click **Next** to continue.
6. Create a target appliance by specifying the following options.

- **Name**—Specify the virtual machine display name for the target appliance. Because the guest name cannot exceed 15 characters, this field will be limited to 15 characters if you have **Same as name** enabled.
- **Host Name**—Specify the guest name for the target appliance. If you want it to be the same as the display name, click **Same as name**. The guest name cannot exceed 15 characters.
- **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. You can select **Specify** and identify the amount of memory and the number of cores per socket for the target appliance. 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.
- **Network**—Select the network that you want the target appliance to use. If you do not know your available networks, check with Carbonite.
- **Adapter**—You can select the E1000, E1000E, or VMXNET 3 adapter type.
- **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.

7. 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.

The target appliance will automatically power on during creation.

It may take a while for the target appliance to finish creation, for example 10-20 minutes. Target appliance creation time is dependent on the Carbonite hardware and your connection to the cloud.

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.

Chapter 7 **Jobs**

One you have your source servers in your servers list, you can protect those servers. 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.

On the **Jobs** page, you will find high-level information and controls for your jobs. You can sort the job list by clicking on any column heading.

You can select multiple groups on this page by selecting the checkbox to the left of a group table row. A checkmark indicates the group is selected. You can also select the checkbox column heading to quickly select or deselect all groups.

You can see 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 shown in the **Source** and **Target** table headings will change depending on where you are in your job lifecyle.

- **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 failback**—During the restoring and failback states, the source of the job is your replica server in the cloud and the target of the job is your failback source.

- **Toolbar and table controls**—The following controls are available on the toolbar on the **Jobs** page.
  - **Add**—Click this button to create a new job. See *Protecting servers* on page 29 for more details on this process.
  - **Delete**—Click this button to delete the selected job groups.


- **Select All**—Click this link to select all items in the table. This option will select all table rows across all pages, even if hidden by a search filter. If you want to select only the rows that are visible on the current page, select the checkbox column heading.

- **Clear All Selections**—Click this link to deselect all items in the table. This option will deselect all table rows across all pages, even if hidden by a search filter. If you want to deselect only the rows that are visible on the current page, deselect the checkbox column heading.

- **per page**—Select the number of table rows to display per page.

- **Starting typing to filter**—Text entered in the filter box will narrow the list displayed to only those rows that contain the search text.

- **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.

- **Hyperlink control**—Click the name of a job group to see details for that group. See *Viewing job details* on page 53.

- **Job status**—The **Status** column has different colors and status information to indicate the health of your jobs.
    - **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.

- **Action menu controls**—The following menu options are available in the **Actions** menu for job groups or the individual servers within 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.
    - **Snapshot**—Select this option to take a snapshot.
    - **Failover**—Select this option to begin failover. See *Failing over servers* on page 39 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 server in the cloud and restores it back to your original source or another server. See *Restoring servers* on page 46 for more details on this process.
    - **Failback**—Select this option after restoration is complete, to finalize the failback process. See *Failback* on page 52 for more details on this process.
    - **Start**—Select this option to start a paused or stopped job. You can also use this action to restart a failed operation.
    - **Stop**—Select this option to stop the job. Data changes will not queue on the source (if you are protecting) or replica server (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 server (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 server created in the cloud will be deleted and the job will be restarted.
- **Reprotect**—Select this option to restart 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 server created in the cloud from the last job will be deleted.
- **Delete**—Select this option to delete the job.
- **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 servers (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 servers to the failback sources (when restoring).
    - **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 servers to the failback sources (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 servers and failback sources (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 source compared to the replica server they are restoring from. For protection, this is the longest time period of data that would be lost if a failure were to occur at the current time.
    - **Replication Queue**—This is the amount of disk space being used to queue replication data on the source servers (when protecting) or replica servers (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 servers to the failback sources (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 servers to the failback sources (when restoring). If compression is disabled, this statistic will be the same as bytes sent.

Protecting servers

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.

1. There are multiple ways to begin the protection process.
   - **Jobs page**—On the **Jobs** page, click **Add**.
   - **Servers page**—On the **Servers** page, you have two options.
       - Select one or more unprotected source servers and click **Protect** from the toolbar
       - Select **Protect** from an unprotected source server's **Actions** menu.
2. On the **Servers** section of the protection wizard, verify the servers you want to protect are selected.
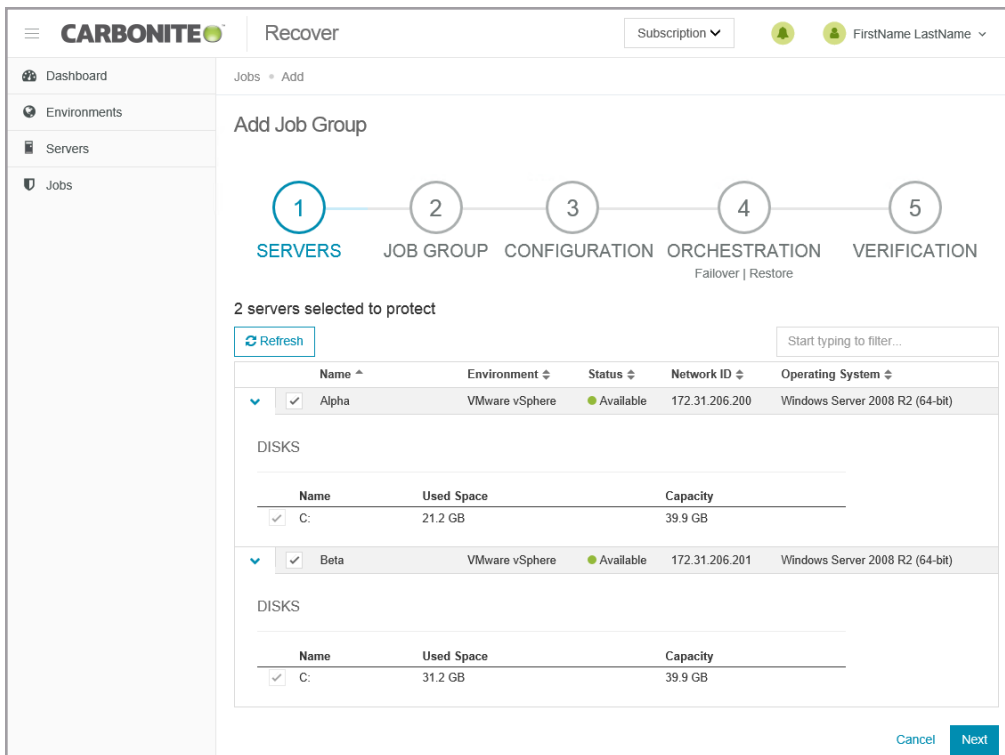
3. Click **Next** to continue.
4. On the **Job Group** section of the protection wizard, specify the settings for the job group.
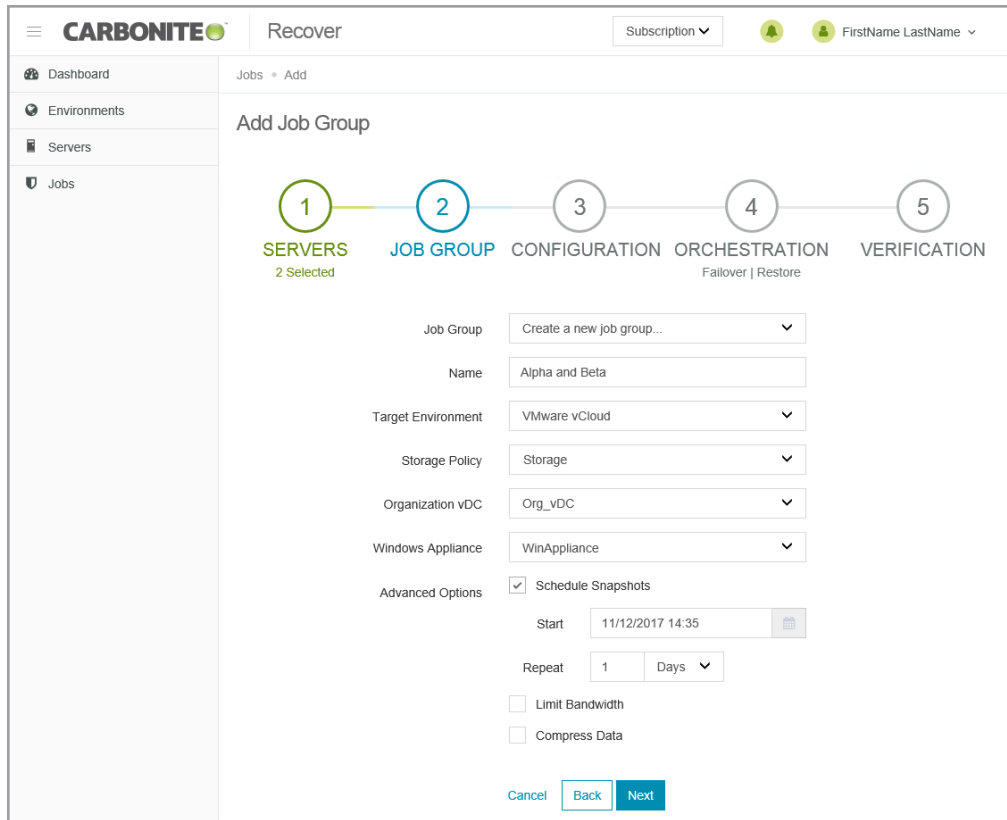
- **Job Group**—At this time, you can only select to create a new job group.
- **Name**—Specify a unique name for the group. If that group name already exists, Carbonite Recover will append a unique number to the group name.
- **Target Environment**—Select the cloud environment where you want to protect 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. You can customize it for each server you are protecting in the next step of the protection wizard.
- **Organization vDC**—Select the organization vDC to use from the storage policy selected. If you do not know your organization vDC, check with Carbonite. This is a default selection. You can customize it for each server you are protecting in the next step of the protection wizard.
- **Windows Appliance**—Select a target appliance to use. This is a default group selection. 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 23.
- **Advanced Options**—The advanced settings are job group level settings. All jobs in the group will have the same scheduled snapshots, bandwidth, and compression settings.
    - **Schedule snapshots**—A snapshot is an image of the job group taken at a single point in time. You can failover to a snapshot. However, you cannot access the snapshot to recover specific files or folders. Select this option if you want Carbonite Recover to take snapshots of the job group automatically at set intervals. Specify

when you want to start taking snapshots and how often you want them taken. Make sure you have reviewed the *Requirements* on page 7 for using snapshots.

- **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. Select this option if you want to limit bandwidth usage. Specify the limit amount. Carbonite Recover will not exceed that amount. If you do not select this option, Carbonite Recover will use 100% bandwidth availability.

    > The minimum bandwidth limit is .028 megabits which is 3500 bytes.

- **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. 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. On the **Configuration** section of the protection wizard, specify the settings for each individual server you are protecting.
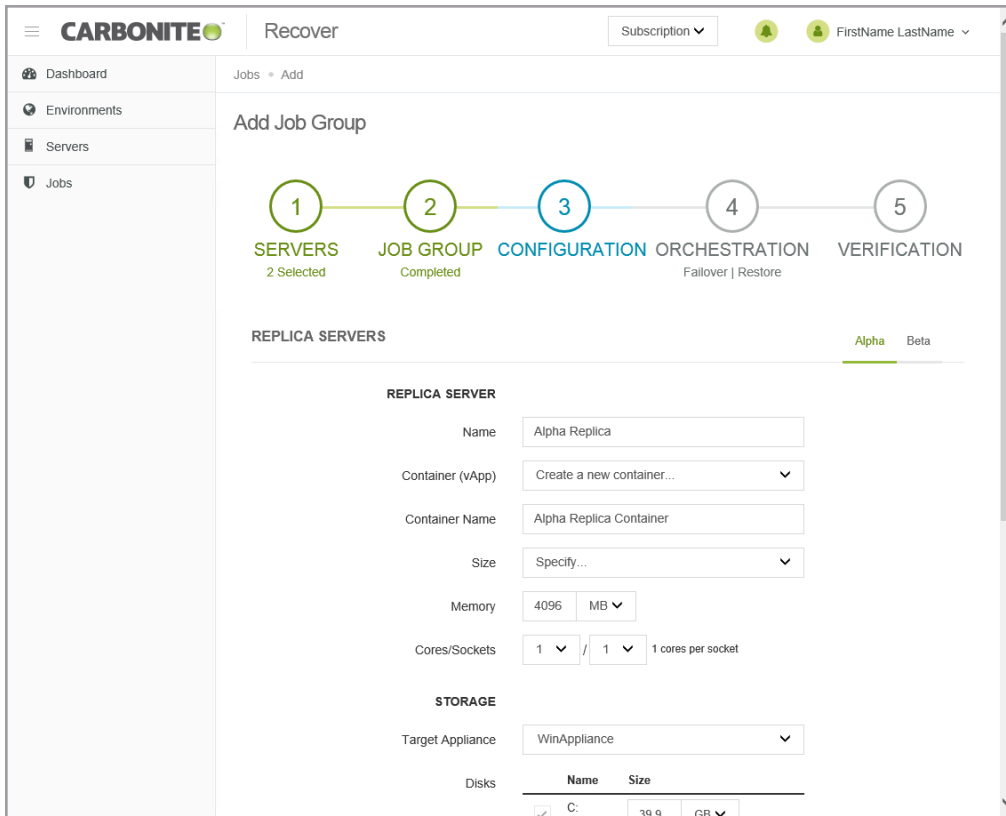
- **Replica Servers**—Specify how you want the virtual machine to be created in the cloud during failover.
    - **Name**—Specify the virtual machine display name. This is the replica server that will be created in the cloud. By default, this is the name of the original source server with the suffix Replica.
    - **Container (vApp)**—Specify if you want to use an existing container in the cloud or create a new one.
    - **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 replica server. You can select **Specify** and identify the amount of **Memory** and the number of **Cores/Socket** for the replica server. 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 server.
    - **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.
    - **Disks**—For each volume on the source server, specify how large you want the corresponding volume to be on the replica server. 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.

- **Network**—For each network adapter on the source server, specify how you want the corresponding adapter to be configured on the replica server. Click **Edit** to change any of the fields for that adapter. After you have made the changes, click **Save**.

    - **Network**—Select the network that you want the adapter to use on the replica server. If you do not know your available networks, check with Carbonite.
    - **Adapter Type**—Select the E1000, E1000E, or VMXNET 3 adapter type.

        The VMXNET 3 adapter is only supported if you have VMware Tools installed on your source server.

    - **IP Mode**—Select **Pool** if you want the replica server to be assigned an IP address from a pool of addresses. Select **Manual** if you want to assign a specific IP address to the adapter. If you select manual, you will have the following options.
        - **IP Addresses**—Specify the IP address you want assigned to the adapter. Click **+** to add another row to the table or **-** to remove an existing row from the table.
        - **Subnet**—Specify the subnet mask to apply to the adapter.
        - **Gateway**—Specify the gateway to apply to the adapter.
        - **DNS Addresses**— Specify the DNS addresses to apply to the adapter. 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.

    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. 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. Set your failover settings, then click **Next** and set your restore and failback settings.

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

    - **Add Failover Plan Details**—This is the first section you will see when you get to step 4. After you have configured the failover plan, click **Next** to see the **Add Restore Plan Details**.

- **Pre-Failover Script**—Before failover starts, you can have your own script launched on a particular server. This script must be a PowerShell script.

  - **Script Name**—Browse (by default the local machine) and select the PowerShell 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. The argument list must be a space separated list and follow standard PowerShell usage rules for string quoting and using special characters.

  - **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, 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.

> 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.

- **Server Failover Order**—If you have more than one server in your group, you can set the failover order of the servers. Drag and drop the servers in the group to the order you want them failed over. Servers in the list will not power on until the replica server before it in the startup order is online. Online in this context means the underlying Recover replication agent service is available for communication.

  > You will have the opportunity to rearrange or disable the startup order before the failover process is started, if desired.
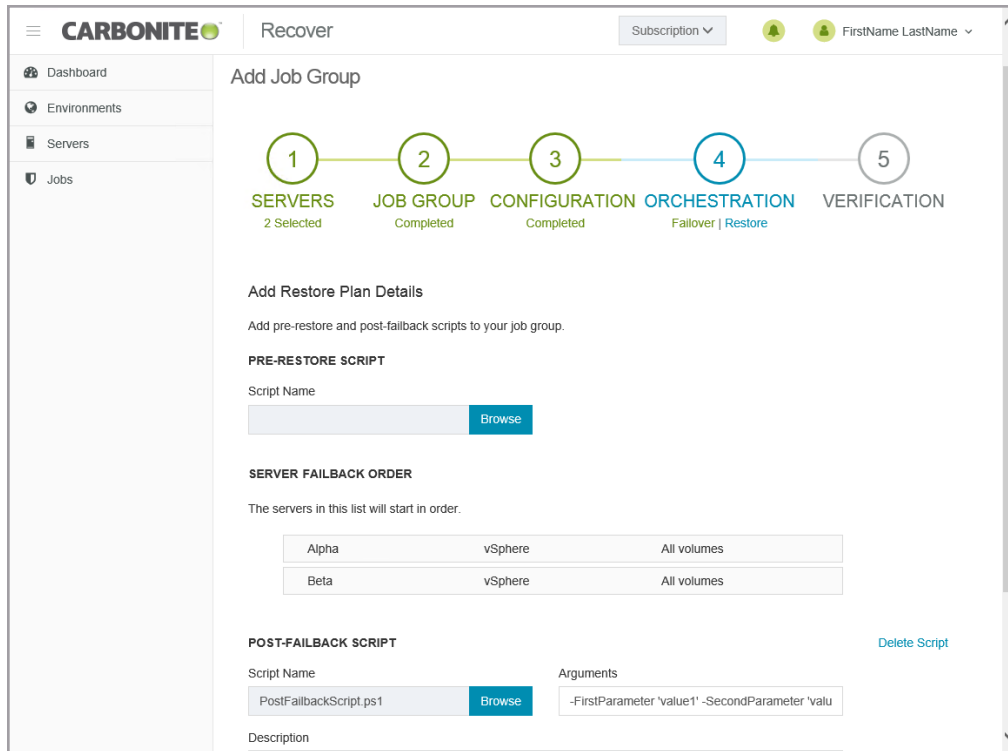
- **Post-Failover Script**—After the failover process is completed (when the last server in the startup order is online), you can have your own script launched on a particular server. This script must be a PowerShell script.
  - **Script Name**—Browse (by default the local machine) and select the PowerShell 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. The argument list must be a space separated list and follow standard PowerShell usage rules for string quoting and using special characters.
  - **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.

  > 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.

- **Add Restore Plan Details**—This is the second section you will see when you get to step 4, after you have clicked **Next** once.

- **Pre-Restore Script**—Before failback starts, you can have your own script launched on a particular server. This script must be a PowerShell script. Keep in mind, this script actually runs before failback starts, not before the restore as currently named.

  - **Script Name**—Browse (by default the local machine) and select the PowerShell 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. The argument list must be a space separated list and follow standard PowerShell usage rules for string quoting and using special characters.

  - **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.

  - **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.

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.

- **Server Failback Order**—If you have more than one server in your group, you can set the failback order of the servers. 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 replica server before it in the startup order has completed failback. Online in this context means the underlying Recover replication agent service is available for communication.

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

- **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 script.
  - **Script Name**—Browse (by default the local machine) and select the PowerShell 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. The argument list must be a space separated list and follow standard PowerShell usage rules for string quoting and using special characters.
  - **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.

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.

9. Click **Next** to continue.
10. Carbonite Recover validates that each source server and target appliance is compatible, as well as validating any group settings. The **Verification** page displays the validation items. Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point

(!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. You must correct any errors before you can continue.

Expand a source server name to see the validation items associated with that server. You can revalidate by clicking **Recheck** or search for specific items by using the filter.

11. Once your configuration has passed verification, click **Finish** to start protection.
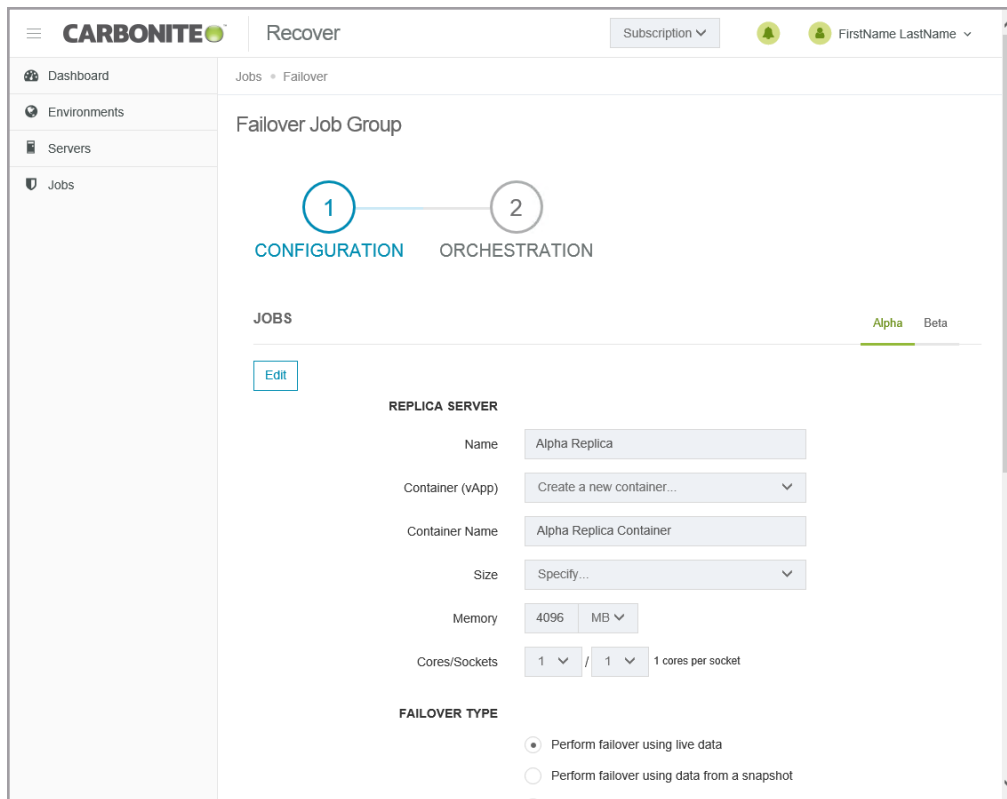
Failing over servers

In the event you need to failover one or more source servers, you can quickly fail them over in the cloud. The failover process workflow is different if you are failing over at the group level or at an individual level.

- **Group level**—Use the group level instructions if you are failing over a group of multiple severs 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 severs or if you are failing over a group that has 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 **Actions** 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**.

- **Replica Server**—Specify how you want the replica server to be created in the cloud during failover.

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

- **Container (vApp)**—Specify if you want to use an existing container in the cloud or create a new one.

- **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 replica server. You can select **Specify** and identify the amount of **Memory** and the number of **Cores/Socket** for the replica server. 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 snapshot**—If you have taken snapshots, you have the option of failing over data from a snapshot. When you select this option, the list of available snapshots will appear. Select the snapshot you want to failover to. The replica data on the target appliance will be reverted to that snapshot and then failover will be initiated. The **Status** and **Description** help you understand what snapshots are available.
  - **Perform test failover**—Select this option to perform a test failover. The test will use the current data on the target appliance, leave the source server online, stop the job, and start the replica server in the cloud using the network settings you configured for the replica server. When the test is complete, you can access the replica server in the cloud to complete your testing. When you are finished with your testing, you can undo the test which will delete the replica server in the cloud and restart the job.

    When failing over a group, you cannot mix test failover with live or snapshot failover within a group. Test failover must be performed for all servers in the group. You can mix live and snapshot failover within a group.

- **Networks**—For each network adapter on the source server, specify how you want the corresponding adapter to be configured on the replica server. Click **Edit** to change any of the fields for that adapter. After you have made the changes, click **Save**.
  - **Network**—Select the network that you want the adapter to use on the replica server. If you do not know your available networks, check with Carbonite.
  - **Adapter Type**—Select the E1000, E1000E, or VMXNET 3 adapter type.

    The VMXNET 3 adapter is only supported if you have VMware Tools installed on your source server.

- **IP Mode**—Select **Pool** if you want the replica server to be assigned an IP address from a pool of addresses. Select **Manual** if you want to assign a specific IP address to the adapter. If you select manual, you will have the following options.
  - **IP Addresses**—Specify the IP address you want assigned to the adapter. Click **+** to add another row to the table or **-** to remove an existing row from the table.
  - **Subnet**—Specify the subnet mask to apply to the adapter.
  - **Gateway**—Specify the gateway to apply to the adapter.
  - **DNS Addresses**— Specify the DNS addresses to apply to the adapter. 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.
3. Click **Next** to continue.
4. Review and if needed, modify your failover plan.
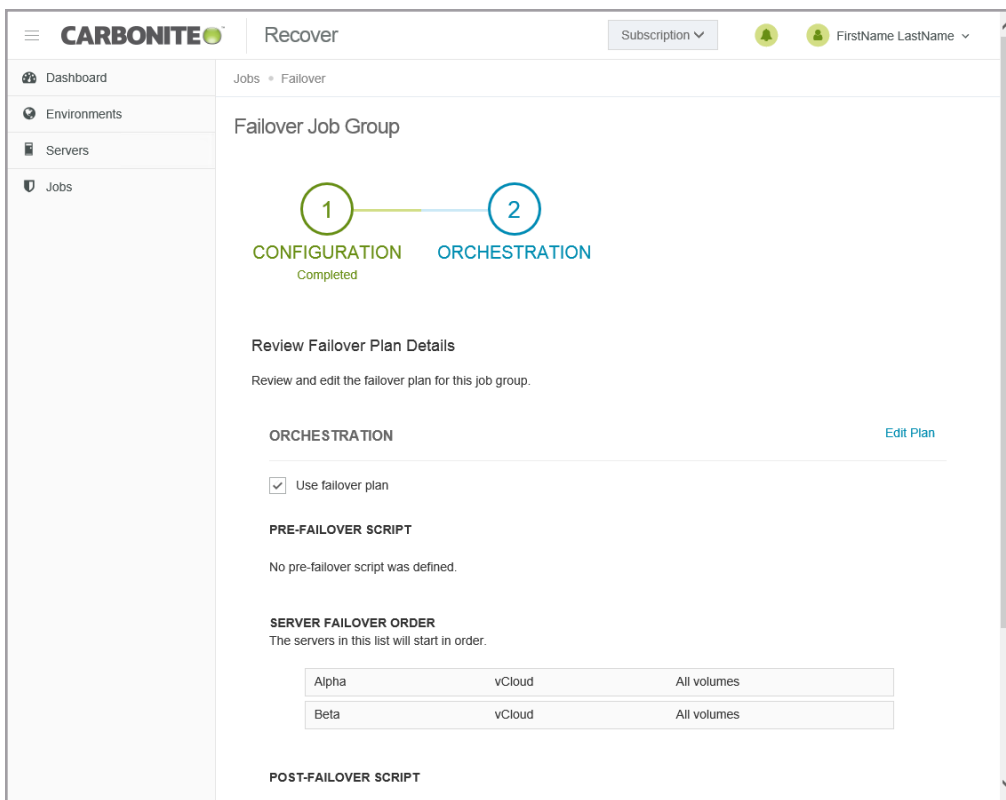
- **Use failover plan**—Enable this option to start the servers in the order specified. You have the option of enabling or disabling scripts as desired. If you disable this option, the servers will all start at the same time and the scripts will be automatically disabled.
- **Disable scripts**—If you are using the failover plan, you can disable scripts so they do not run. If you are not using the failover plan, scripts will automatically be disabled.
- **Enable scripts**—If you are using the failover plan and you have disabled scripts, you can enable them with this option. If you are not using the failover plan, scripts will automatically be 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 script.
    - **Script Name**—Browse (by default the local machine) and select the PowerShell 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. The argument list must be a space separated list and follow standard PowerShell usage rules for string quoting and using special characters.
    - **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, 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.

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

- **Server Failover Order**— You can set the failover order of the servers. Drag and drop the servers in the group to the order you want them failed over. Servers in the list will not power on until the replica server before it in the startup order is online. Online in this context means the underlying Recover replication agent service is available for communication.

- **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 script.
    - **Script Name**—Browse (by default the local machine) and select the PowerShell 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. The argument list must be a space separated list and follow standard PowerShell usage rules for string quoting and using special characters.
    - **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.

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

5. Click **Finish** to start the failover.

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 severs 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 **Actions** 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**.

> If you entered failover scripts during job creation, 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 a single server and not the group.

- **Replica Server**—Specify how you want the replica server to be created in the cloud during failover.
  - **Name**—Specify the virtual machine display name. This is the replica server that will be created in the cloud. By default, this is the name of the original source server with the suffix Replica.
  - **Container (vApp)**—Specify if you want to use an existing container in the cloud or create a new one.
  - **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 replica server. You can select **Specify** and identify the amount of **Memory** and the number of **Cores/Socket** for the replica server. 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 snapshot**—If you have taken snapshots, you have the option of failing over data from a snapshot. When you select this option, the list of available snapshots will appear. Select the snapshot you want to failover to. The replica data on the target appliance will be reverted to that snapshot and then failover will be initiated. The **Status** and **Description** help you understand what snapshots are available.
  - **Perform test failover**—Select this option to perform a test failover. The test will use the current data on the target appliance, leave the source server online, stop the job, and start the replica server in the cloud using the network settings you configured for the replica server. When the test is complete, you can access the replica server in the cloud to complete your testing. When you are finished with your testing, you can undo the test which will delete the replica server in the cloud and restart the job.

- **Networks**—For each network adapter on the source server, specify how you want the corresponding adapter to be configured on the replica server. Click **Edit** to change any of the fields for that adapter. After you have made the changes, click **Save**.
  - **Network**—Select the network that you want the adapter to use on the replica server. If you do not know your available networks, check with Carbonite.
  - **Adapter Type**—Select the E1000, E1000E, or VMXNET 3 adapter type.

    > The VMXNET 3 adapter is only supported if you have VMware Tools installed on your source server.

  - **IP Mode**—Select **Pool** if you want the replica server to be assigned an IP address from a pool of addresses. Select **Manual** if you want to assign a specific IP address to the adapter. If you select manual, you will have the following options.
    - **IP Addresses**—Specify the IP address you want assigned to the adapter. Click **+** to add another row to the table or **-** to remove an existing row from the table.
    - **Subnet**—Specify the subnet mask to apply to the adapter.
    - **Gateway**—Specify the gateway to apply to the adapter.
    - **DNS Addresses**— Specify the DNS addresses to apply to the adapter. 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.

3. Click **Failover** to start the failover.

Restoring servers

After you have failed over to the cloud, you can restore from the replica server 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.

- **Group level**—Use the group level instructions if you are restoring a group of multiple severs 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.
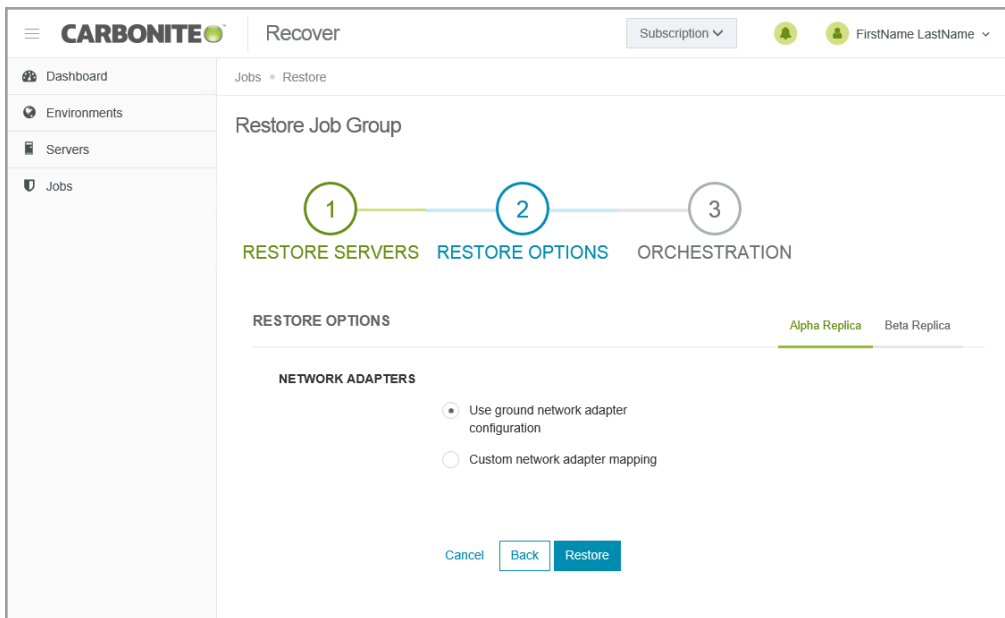
Restoring at the group level

Use this process if you are restoring a group of multiple severs 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 server uses same IP address as original source**—If you chose to use the same IP address on the replica server in the cloud as was used on any original sources, you must bring the original sources up offline, assign them a new IP address, and then bring them online. If you are restoring to any different servers, make sure they have a unique IP address when they are brought online.
    - **Replica server uses different IP address than original source**—If you chose to use a different IP address on the replica server in the cloud than was used on any original sources, no additional steps are required. You can bring the original sources online as is or use different servers for restoration.
2. On the **Jobs** page, select **Restore** from the **Actions** 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 server you are restoring, select the failback source 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 server which is now standing in for your original source. Your selected servers must be online before you can continue.
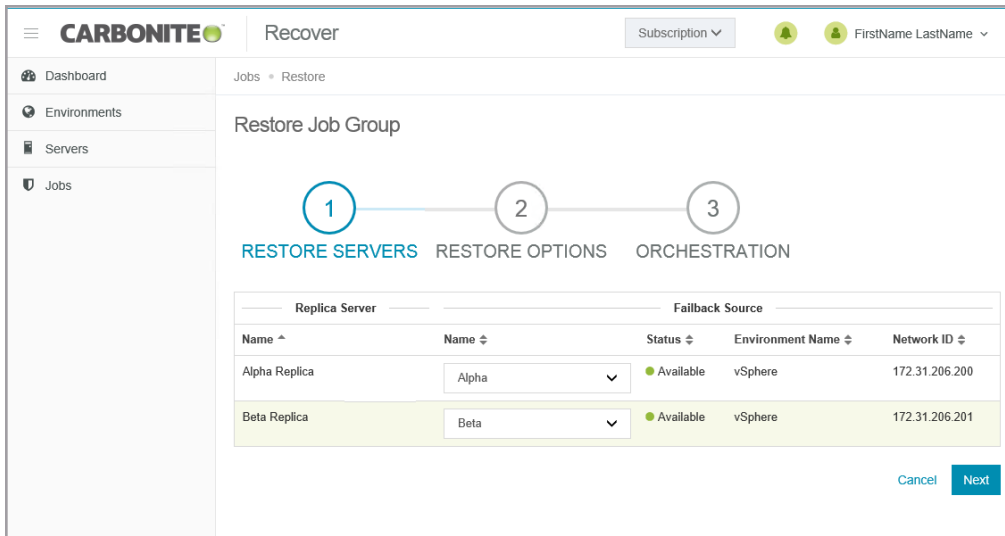
4. Click **Next** to continue.
5. For each replica server you are restoring, specify how you want to handle the network adapters. Specify the network adapters for each replica server by selecting the server tabs.

- **Use ground network adapter configuration**—This option will leave the configuration of the network adapters on the failback source 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 server to the failback source. Map the network adapters from your replica server to the adapters on the failback source server. You can also choose to ignore the network adapters from your replica server. Ignoring the network adapter will not use it on the failback source. Any network adapters on the failback source server that are not mapped to a replica server adapter will be left as is.

6. Click **Next** to continue.
7. Review and if needed, modify your restore plan.
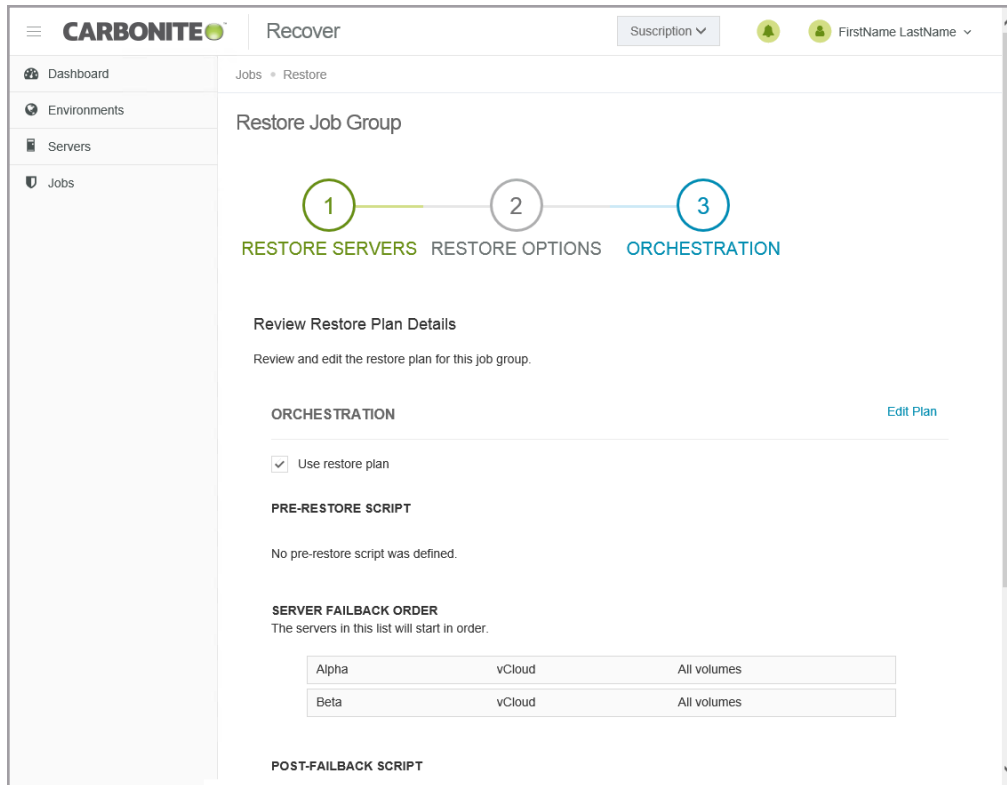
- **Use restore plan**—Enable this option to start the servers during failback in the order specified. You have the option of enabling or disabling scripts as desired. If you disable this option, the servers will all start at the same time during failback and the scripts will be automatically disabled.

- **Disable scripts**—If you are using the restore plan, you can disable scripts so they do not run. If you are not using the restore plan, scripts will automatically be disabled.

- **Enable scripts**—If you are using the restore plan and you have disabled scripts, you can enable them with this option. If you are not using the restore plan, scripts will automatically be 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-Restore Script**—Before failback starts, you can have your own script launched on a particular server. This script must be a PowerShell script. Keep in mind, this script actually runs before failback starts, not before the restore as currently named.

    - **Script Name**—Browse (by default the local machine) and select the PowerShell 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. The argument list must be a space separated list and follow standard PowerShell usage rules for string quoting and using special characters.

    - **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.
- **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.

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

- **Server Failback Order**— You can set the failback order of the servers. 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 replica server before it in the startup order has completed failback. Online in this context means the underlying Recover replication agent service is available for communication.

- **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 script.
  - **Script Name**—Browse (by default the local machine) and select the PowerShell 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. The argument list must be a space separated list and follow standard PowerShell usage rules for string quoting and using special characters.
  - **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.

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

8. Click **Restore** 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 severs or if you are restoring a group that has only one server, use the restoring at the group level instructions.
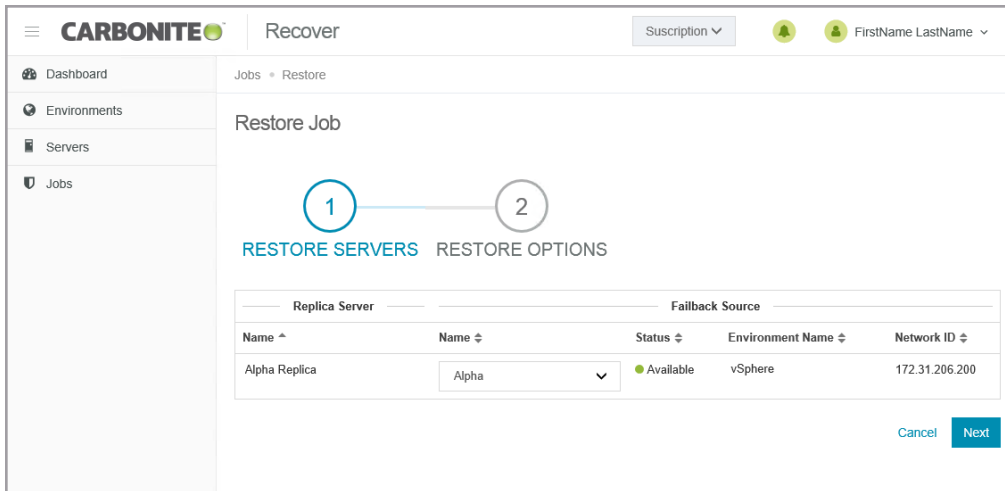
1. Your first step will depend on how you failed over your original source to the cloud.
   - **Replica server uses same IP address as original source**—If you chose to use the same IP address on the replica server in the cloud as was used on the original source, you must bring the original source up 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 server uses different IP address than original source**—If you chose to use a different IP address on the replica server in the cloud than was used on the original source, no additional steps are required. You can bring the original source online as is or use a different server for restoration.
2. On the **Jobs** page, select **Restore** from the **Actions** menu for an individual server.
3. For the replica server you are restoring, select the failback source 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 server which is now standing in for your original source. Your selected server must be online before you can continue.

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

> If you entered restore or failback scripts during job creation, you will see a warning on the **Restore Job** page. It is a notification that scripts will not be run since you are only restoring a single server and not the group.

- **Use ground network adapter configuration**—This option will leave the configuration of the network adapters on the failback source 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 server to the failback source. Map the network adapters from your replica server to the adapters on the failback source server. You can also choose to ignore the network adapters from your replica server. Ignoring the network adapter will not use it on the failback source. Any network adapters on the failback source server that are not mapped to a replica server adapter will be left as is.

6. Click **Restore** to start the restoration.

Failback

After the restoration is complete, you need to complete failback. This finalizes the identity transfer from the replica server to the failback source.

On the **Jobs** page, select **Failback** from the **Actions** 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 source will be a replica of your replica server in the cloud, including any changes that were made to that replica server while it was running in the cloud.

You can reprotect the server again by selecting **Reprotect** from the **Actions** menu. If you have not deleted the replica server in the cloud, you can reuse the hard disks that were already created. You can also create new disks, if desired.

Viewing job details

On the **Jobs** page, click the name of a group to see job details. On the **Details** page, you will find group statistics as well as individual server details and statistics.

- **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 servers (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 servers to the failback sources (when restoring).

  - **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 servers to the failback sources (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 servers and failback sources (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 source compared to the replica server they are restoring from. For protection, this is the longest time period of data that would be lost if a failure were to occur at the current time.

  - **Replication Queue**—This is the amount of disk space being used to queue replication

data on the source servers (when protecting) or replica servers (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 servers to the failback sources (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 servers to the failback sources (when restoring). If compression is disabled, this statistic will be the same as bytes sent.

- **Job Details**—Click a server name tab to see server details, operations, statistics, and snapshots for that individual server.

  - **Source Server**—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.

  - **Target Appliance**—During the restoring and failback states, the source of the job is your replica server in the cloud and the target of the job is your failback source.

  - **Status**—The **Status** detail has different colors and status information to indicate the health of your jobs.

    - **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.

- **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.

- **Statistics**—This section shows the statistics for the individual server.

  - **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 servers (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 servers to the failback sources (when restoring).

  - **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 servers to the failback sources (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 servers and failback sources (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 source compared to the replica server they are restoring from. For protection, this is the longest time period of data that would be lost if a failure were to occur at the current time.

- **Replication Queue**—This is the amount of disk space being used to queue replication data on the source servers (when protecting) or replica servers (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 servers to the failback sources (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 servers to the failback sources (when restoring). If compression is disabled, this statistic will be the same as bytes sent.

- **Snapshots**—If there are any snapshots for the server, they will be listed. You can use snapshots to failover to an earlier point in time. To help you understand what snapshots are available, the **Description** indicates if the snapshot was scheduled or manually taken (taken on demand). If you no longer need a snapshot, click **Delete** to remove it from the job.

Chapter 8 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 your sign in name, in the upper right corner of the Carbonite Recover web page, and select **User Preferences**.

2. In the **Email** section, modify your email prefix and notification 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. This prefix allows you to recognize and filter emails specific to Carbonite Recover. You can change or remove the prefix as desired. The remainder of the subject line will contain the 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**.

Chapter 9 **Subscription usage**

You can review your subscription usage at any time.

1. Usage billing is per subscription. If you are assigned to multiple subscriptions, select the subscription that you want to view subscription usage for by selecting the subscription name in the drop-down list next to the bell notification icon.

2. Click on your sign in name, in the upper right corner of the Carbonite Recover web page and select **Subscription Usage**.

> It may take a minute for the usage information to populate on screen.

3. Select a time period from the drop-down list in the upper right corner of the **Subscription Usage** page. 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.

4. Review the information provided in the charts and tables.
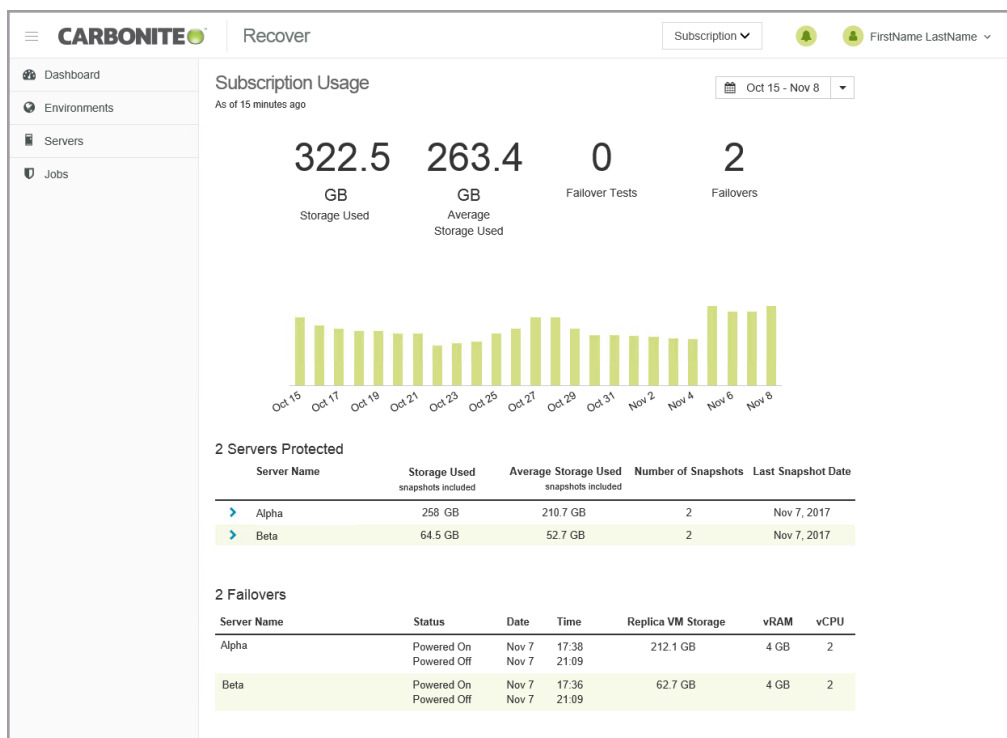
- **High level statistics**—At the top of the page, you will find high level statics for your subscription.
    - **Storage Used**—This is the amount of storage used for the last day of the selected billing period.
    - **Average Storage Used**—This is the average amount of storage used for the selected 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 snapshot 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 snapshots for the server and the date of the last snapshot. 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, snapshot, or test). You can see when the replica server was powered on and off along with statistics for the replica server.