


## *User's Guide*



## Notices

Carbonite Migrate User's Guide, Wednesday, October 27, 2021

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

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

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

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

© 2021 Carbonite, Inc. All rights reserved.

# Contents

---

<b>Chapter 1 Carbonite Migrate overview</b>	<b>5</b>
Configurations and ports	6
<b>Chapter 2 Requirements</b>	<b>10</b>
Source requirements	11
Target requirements	15
Target requirements for Amazon Web Services	16
Target requirements for Google Cloud Platform	17
Target requirements for Microsoft Azure Resource Manager	18
Registering and assigning owner role in Microsoft Azure Resource Manager	19
Target requirements for Microsoft Azure Resource Manager (legacy)	24
Target requirements for Microsoft Azure Classic	25
Target requirements for custom environments	26
Proxy requirements	29
Carbonite Migrate interface requirements	30
<b>Chapter 3 Getting started</b>	<b>31</b>
<b>Chapter 4 Carbonite Migrate interface</b>	<b>32</b>
Configuration	34
<b>Chapter 5 Environments</b>	<b>36</b>
Creating an Amazon Web Services environment	39
Minimum permissions for Amazon Web Services	41
Creating a Google Cloud Platform environment	43
Creating a Microsoft Azure Resource Manager environment	45
Creating a Microsoft Azure Resource Manager (legacy) environment	47
Creating a Microsoft Azure Classic environment	49
Creating a VMware vCloud Director environment	51
Creating a VMware vSphere environment	53
Creating an Active Directory environment	54
Creating a custom environment	56
Proxy server	57
<b>Chapter 6 Servers</b>	<b>58</b>
Discovering servers in non-custom environments	60
Discovering servers in custom environments	61
Viewing server details	63
<b>Chapter 7 Migrations</b>	<b>64</b>
Migrating to Amazon Web Services	67
Migrating to Google Cloud Platform	73
Migrating to Microsoft Azure Resource Manager	79
Migrating to Microsoft Azure Resource Manager (legacy)	86
Migrating to Microsoft Azure Classic	92
Migrating to a custom environment	98
Viewing migration details	102
Viewing migration history	103

---

<b>Chapter 8 Email notification .....</b>	<b>104</b>
<b>Chapter 9 Account management .....</b>	<b>105</b>

---

## Chapter 1 Carbonite Migrate overview

Carbonite Migrate is an online service that enables migrations of Windows and Linux servers to a public or private cloud. Using only a few clicks in Carbonite Migrate, you can easily automate server migrations for a secure, near-zero downtime experience.

There are four main components to Carbonite Migrate.

- **Source**—The server you want to migrate is called your source server. The source server can be a physical or virtual server on the ground or in a cloud.
- **Target**—The new server you are migrating to is called the target server. The target server is a virtual server and it can be hosted on the ground or in a cloud.
- **Carbonite Migrate**—This is the online service that enables the migration of the source to the target.
- **Proxy**—If your source or target does not have a public IP address, you will have to use a proxy server for communication between the source or target and the Carbonite Migrate service. The proxy server is not used for communication between the source and target.

See *Configurations and ports* on page 6 for details and diagrams of how to configure Carbonite Migrate.



Due to limitations provisioning storage for Linux servers in the cloud, Linux source servers can only be migrated to custom target environments at this time. This means you will need to have an existing target server already provisioned and running in your cloud to use as a custom target.

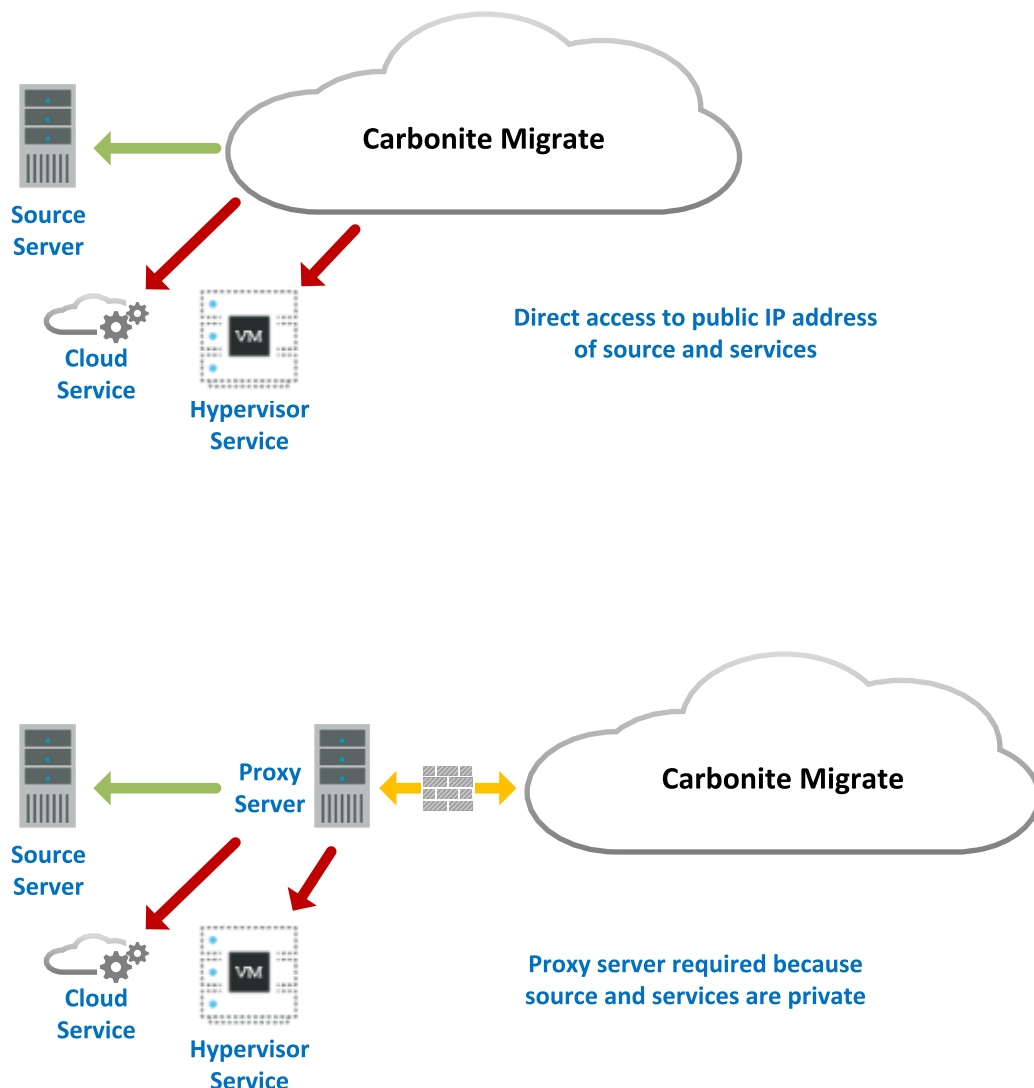
---

# Configurations and ports

The flexibility of private and public accessibility allows for many different Carbonite Migrate configurations, including whether a proxy server is needed. If your source or target does not have a public IP address, you will have to use a proxy server for communication between the source or target and the Carbonite Migrate service. The proxy server is not used for communication between the source and target. Even though migration data is encrypted, you may need to consider private and public accessibility between your source and target for the migration of the data. There are different port requirements depending on the configuration you are using.

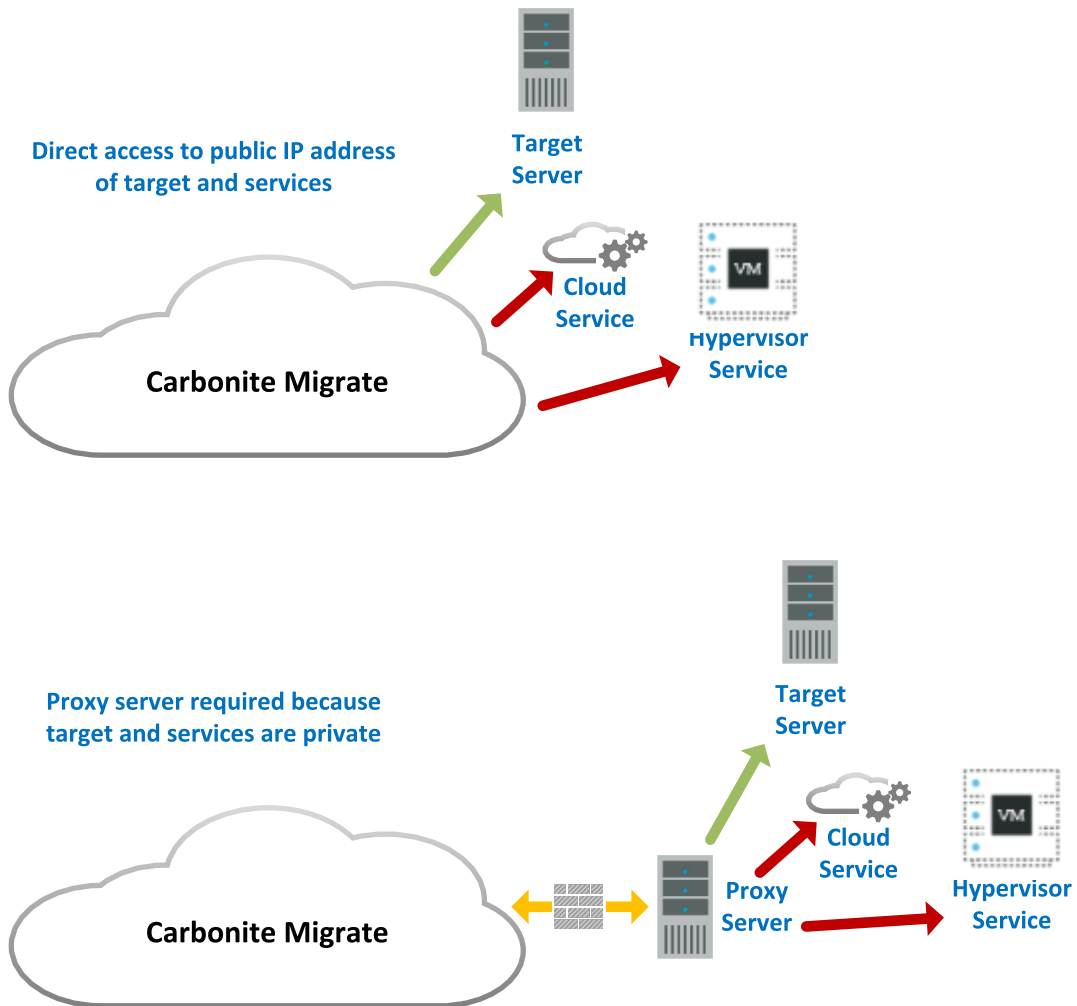
## Source configurations

Your source environment will have the source server and optionally cloud services or a hypervisor, depending on your configuration.



## Target configurations

Your target environment will have the target server and cloud services or a hypervisor, depending on your configuration.



## Source to target configurations

Even though migration data is encrypted, you may need to consider private and public accessibility between your source and target for the migration of the data.

Public IP addresses allow the source and target to migrate the encrypted data directly

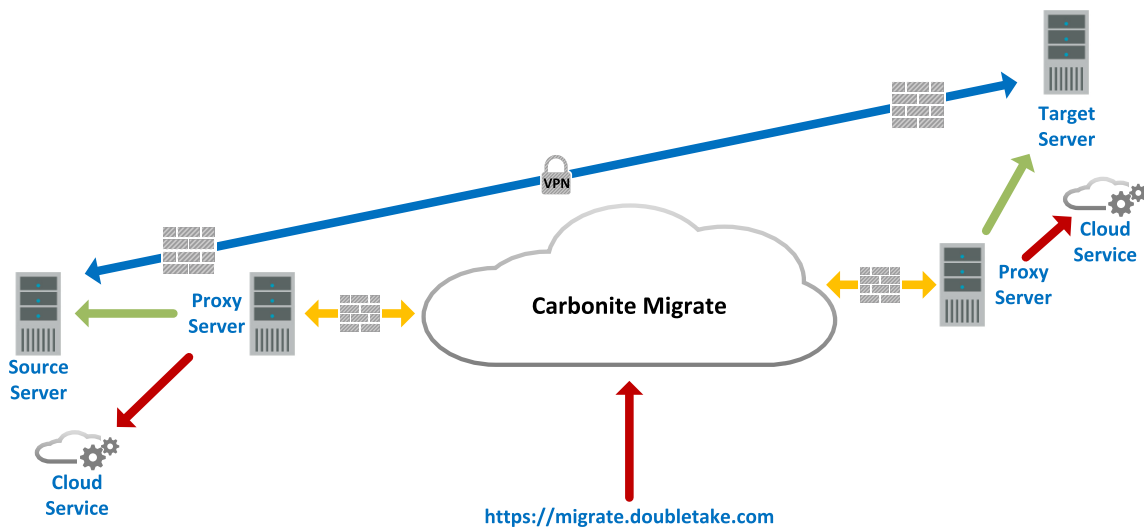


Private IP addresses require a VPN between the source and target to migrate the encrypted data



## Sample configuration





Any combination of the configurations above is possible. For example, if you are migrating from cloud to cloud and your source and target both have private IP addresses, your configuration would look similar to this.





## Ports

The colored arrows correspond to the configuration diagrams above.

Component to Component	Windows Communication	Linux Communication	Arrow Color
Web Browser to Carbonite Migrate Carbonite Migrate to Cloud or Hypervisor Service Proxy Server to Cloud or Hypervisor Service	HTTPS port 443		
Proxy to Carbonite Migrate	Service Bus (outbound only) port 9354 SignalR port 31867		
Proxy to Source Proxy to Target Carbonite Migrate to Source Carbonite Migrate to Target	SMB port 445 WinRM ports 5985 and 5986 Carbonite Migrate port 6325	SSH port 22 Carbonite Migrate port 6326	
Source to Target	Carbonite Migrate ports 6320 and 6325	Carbonite Migrate ports 1500, 6325, and 6326	

---

## Chapter 2 Requirements

Review the separate requirements sections for details on the requirements for your source servers and source environments, your target servers and target environments, your proxy server if you are using one, and for the Carbonite Migrate service and interface.

- *Source requirements* on page 11—The servers you are migrating have specific requirements including operating system and configuration of the source servers. There are also requirements for the environment where the source servers are located, networking, and other specifications.
- *Target requirements* on page 15—In addition to general target requirements, you must meet requirements specific to the target cloud environment you are using.
- *Proxy requirements* on page 29— If your source or target does not have a public IP address, you will have to use a proxy server for communication between the source or target and the Carbonite Migrate service. The proxy server is not used for communication between the source and target.
- *Carbonite Migrate interface requirements* on page 30—Carbonite Migrate is the online service and interface that you use to execute and manage the migration. You must meet the specific requirements to use the service.



Due to limitations provisioning storage for Linux servers in the cloud, Linux source servers can only be migrated to custom target environments at this time. This means you will need to have an existing target server already provisioned and running in your cloud to use as a custom target.

---

# Source requirements

The source servers you are migrating, and the environment where the source servers are located, must meet the following requirements.

- **Environment**—The source servers must be running in one of the following supported environments.
  - **Amazon Web Services (AWS)**—You must have an access key and secret key in order to use Amazon Web Services.
  - **Google Cloud Platform (GCP)**—You must have an existing project in Google Cloud and know the project ID. You also must have a service account associated with the Google Cloud project. The account must have, at a minimum, the **Compute Admin** role. You will need the service account email address and the private key associated with the project in order to create your Google Cloud Platform environment in Carbonite Migrate.
  - **VMware vCloud Director**—You must have a version of vCloud Director that is currently supported by VMware. You must know the vCloud IP address or hostname as well as the organization within vCloud that you have access to and want to use. You must have credentials that can access vCloud.
  - **VMware vSphere**—You must have one of the following versions of vSphere: 5.5, 6.0, 6.5, or 6.7. You must know the vSphere IP address or hostname and have credentials to access it.
  - **Active Directory**—Local networks running Active Directory are supported. Carbonite Migrate will query the Active Directory service in order to identify servers that you can migrate. In a domain environment, Carbonite Migrate (or a proxy server) must be able to resolve the domain name and connect to the domain controller using port 389.
  - **Custom** —This environment type is a manually created environment when your servers are not in any of the other environments. Servers can be discovered using a range of IP addresses or an IP address and subnet mask. The servers must be able to respond to a ping request. You can also manually add servers to the environment by a single IP address or name. If you specify a computer name, Carbonite Migrate (or a proxy server) must be able to resolve the computer name for the IP address.



If your custom server is hosted on Windows Azure, the server must be using DHCP.

---

- **Operating system**—The source servers must be 64-bit architecture and must be one of the following Windows or Linux operating systems, with supported file system and kernel type. Make sure your operating system is also supported on the target cloud you are migrating to.

Operating System	Version	File System	Kernel Type
Windows	2008 R2 Service Pack 1 or later	NTFS	Not applicable
	2012		
	2012 R2		
	2016	NTFS	
	2019	ReFS	
Red Hat Enterprise Linux	6.8 through 6.10	Ext3	Default
Oracle Enterprise Linux	7.5 through 7.7	Ext4	
CentOS	8.0	XFS	
SUSE Linux Enterprise	11.2 through 11.4	Ext3 XFS	Default Xen
	12.2 through 12.4	Ext3 Ext4	Default
	15.0 through 15.1	XFS Btrfs	
Ubuntu	16.04.2 through 16.04.4	Ext2 Ext3	Generic
	18.04.0 through 18.04.1	Ext4 XFS	



A Windows source server cannot be a Hyper-V server. Migration of a Hyper-V server is not supported.



Due to limitations provisioning storage for Linux servers in the cloud, Linux source servers can only be migrated to custom target environments at this time. This means you will need to have an existing target server already provisioned and running in your cloud to use as a custom target.

Btrfs cannot be migrated together with ext4. Btrfs and ext4 can be combined with other file systems but not with each other.

For all Linux operating systems except Ubuntu, the kernel version must match the expected kernel for the specified release version. For example, if `/etc/redhat-release` declares the system to be a Redhat 7.6 system, the kernel that is installed must match that.

Stacking filesystems, like eCryptFS, are not supported.

---

- **Linux packages and services**—Carbonite Migrate will check for and, if necessary, install the required packages and services needed to install and use Carbonite Migrate. This process requires Internet access.
- **SELinux policy**—SELinux is supported on your Linux servers.
- **UEFI, trusted boot, secure boot**—UEFI (Unified Extensible Firmware Interface) is supported on the source, however, trusted boot (tboot), secure boot, or other volume blocking mechanisms are not supported.
- **Linux networking**—The network configuration (traditional, NetworkManager, or Netplan) on the source and target should match. If you have a mix of network types, you may have to configure the networking on the target after cutover.
- **Root and extended partitions**—Linux root partitions and extended partitions are not supported and will not be migrated.
- **System memory**—At least 1 GB of memory is required on the source servers.
- **Disk size**—Due to operating system limitations, the system volume on Windows sources cannot be larger than 2 TB.
- **Server name**—Unicode file system support is included, but the source name must still be in ASCII format. If you have the need to use a server's fully-qualified domain name, your server cannot start with a numeric character because that will be interpreted as an IP address. Additionally, all source servers in your environment must have a unique server name.
- **Source to target communication**—A VPN is required if your source and target cannot communicate together in both directions.
- **Source to Carbonite Migrate communication**—If your source does not have a public IP address, you will have to use a proxy server for source to Carbonite Migrate communication. See *Proxy requirements* on page 29 for details.
- **Windows Remote Management**—On Windows servers, you must enable WinRM, by typing **winrm quickconfig** at a local command prompt on each server and then typing **y** to grant administrative rights remotely to local users.
- **Windows firewall configuration**—On each Windows source, you must have the following ports open.
  - **File and Printer Sharing (SMB-In)**—Port 445 must be open for all profiles.
  - **PowerShell and Windows Remote Management (HTTP-In)**—Port 5985 (for HTTP) and 5986 (for HTTPS) must be open for public and domain profiles. Confirm in Windows Remote Management (HTTP-In) that the **Scope** for **Remote IP address** is set to **Any IP address**.
  - **Carbonite Migrate**—Ports 6320, 6325, and 6326 must be open for UDP and TCP for both inbound and outbound traffic.
- **Linux firewall configuration**—On each Linux source, you must have ports 22, 1500, 1505, 1506, 6325, and 6326 open for both UDP and TCP for both inbound and outbound traffic. You must also have port 1501 open for localhost communication.

- **Carbonite licensing**—Carbonite Migrate requires a license on each source server. Source licenses will automatically be activated when a migration job is started and will automatically expire after 30 or 60 days. If you delete and re-create a migration from that same source, the license will be reused, however you still only have the 30 or 60 day window from when it was first activated to complete the migration. See *Configuration* on page 34 for details on adding licenses to your environment.



Once you have activated a license, you will receive an email notification when the license has seven days left until expiration and again at one day left until expiration.

---

- **Time**—The clock on your source and target servers must be within a few minutes of each other, relative to UTC. Large time skews (more than five minutes) will cause errors during cutover.
- **SAP HANA**—SAP HANA is supported. If you are using an SAP HANA database on your source, Carbonite Migrate will automatically quiesce the database and replicate any final data changes during the cutover process.

# Target requirements

The target servers you are migrating to have general requirements as well as requirements specific for the environment where the target servers are located.

- **Environment requirements**—You must meet requirements specific to the target cloud environment you are using. See the following links for requirements specific to that environment type.
    - *Target requirements for Amazon Web Services* on page 16
    - *Target requirements for Google Cloud Platform* on page 17
    - *Target requirements for Microsoft Azure Resource Manager* on page 18
    - *Target requirements for Microsoft Azure Resource Manager (legacy)* on page 24
    - *Target requirements for Microsoft Azure Classic* on page 25
    - *Target requirements for custom environments* on page 26
- 



Use the Microsoft Azure Resource Manager target type if you have a migrate project under Azure Migrate in the Azure portal. A migrate project is used to store discovery, assessment, and migration metadata reported by the on-premise environment. If you do not have a migrate project, you can learn more about them at <https://azure.microsoft.com>. For migrations to Azure Resource Manager without using Azure Migrate projects, use the Microsoft Azure Resource Manager (legacy) target environment type.

---

- **Target to source communication**—A VPN is required if your source and target cannot communicate together in both directions.
- **Target to Carbonite Migrate communication**—If your target does not have a public IP address, you will have to use a proxy server for target to Carbonite Migrate communication. See *Proxy requirements* on page 29 for details.
- **Carbonite licensing**—You do not have to purchase or activate a Carbonite Migrate license on each target server. Carbonite Migrate will handle the licensing on the target server.
- **Time**—The clock on your source and target servers must be within a few minutes of each other, relative to UTC. Large time skews (more than five minutes) will cause errors during cutover.

## Target requirements for Amazon Web Services

In addition to the general *Target requirements* on page 15, an Amazon Web Services environment must meet the following requirements.

- **Keys**—You must have an access key and secret key in order to use Amazon Web Services.
- **Templates**—During the migration setup, you must select an operating system template for each source operating system you are migrating. This template will be used to create the target server in the cloud.
  - **Amazon Machine Images**—The list of available templates is determined by the AMIs (Amazon Machine Images) that you have available. Carbonite Migrate can only use AMIs that meet the following characteristics.
    - **Image type**—machine
    - **Root device type**—ebs
    - **Virtualization type**—hvm
    - **Block device mapping volume type**— standard or gp2

If you do not have an AMI that meets this criteria, your search list will be empty.



## Target requirements for Google Cloud Platform

In addition to the general *Target requirements* on page 15, a Google Cloud Platform environment must meet the following requirements.

- **Project**—You must have an existing project in Google Cloud and know the project ID.
- **Service account**—You must have a service account associated with the Google Cloud project. The account must have, at a minimum, the **Compute Admin** role. You will need the service account email address to create your Google Cloud Platform environment in Carbonite Migrate.
- **Private key**—You must have the private key associated with the service account.
- **Source server operating system**—Only Windows source servers can be migrated to Google Cloud Platform at this time.
- **Templates**—During the migration setup, you must select an operating system template for each source operating system you are migrating. This template will be used to create the target server in the cloud.

## Target requirements for Microsoft Azure Resource Manager

In addition to the general *Target requirements* on page 15, a Microsoft Azure Resource Manager environment must meet the following requirements.

- **Configuration**—You must register Carbonite Migrate as a consumer of Active Directory and assign it an owner role. During the registration, you will need to obtain a **Client ID**, **Client Secret**, and **Tenant ID** that will be required when creating an Azure Resource Manager environment in Carbonite Migrate. See *Registering and assigning owner role in Microsoft Azure Resource Manager* on page 19 for details.
- **Networking**—A virtual network and subnet must be associated to the region you are migrating to.
- **Subscription**—You must have your Azure Active Directory subscription ID.
- **Migration project**— Use the Microsoft Azure Resource Manager target type if you have a migrate project under Azure Migrate in the Azure portal. A migrate project is used to store discovery, assessment, and migration metadata reported by the on-premise environment. If you do not have a migrate project, you can learn more about them at <https://azure.microsoft.com>. For migrations to Azure Resource Manager without using Azure Migrate projects, use the Microsoft Azure Resource Manager (legacy) target environment type.
- **Templates**—During the migration setup, you must select an operating system template for each source operating system you are migrating. This template will be used to create the target server in the cloud.

## Registering and assigning owner role in Microsoft Azure Resource Manager

When using Microsoft Azure Resource Manager, you must register Carbonite Migrate as a consumer of Active Directory and assign it an owner role. During the registration, you will need to obtain a **Client ID**, **Client Secret**, and **Tenant ID** that will be required when creating an Azure Resource Manager environment in Carbonite Migrate.

Complete the following steps in the new Azure portal.

1. To register Carbonite Migrate as a consumer of Active Directory, create a new application registration. Begin the process using one of the following methods.
  - Search for app registrations in the top search bar and then click **New registration**.
  - Select **Azure Active Directory** from the fly out menu and then under **Create**, select **App registration**. (If **Azure Active Directory** is not listed in **Favorites**, select **All services** then **Azure Active Directory**.)
2. Specify the application information.

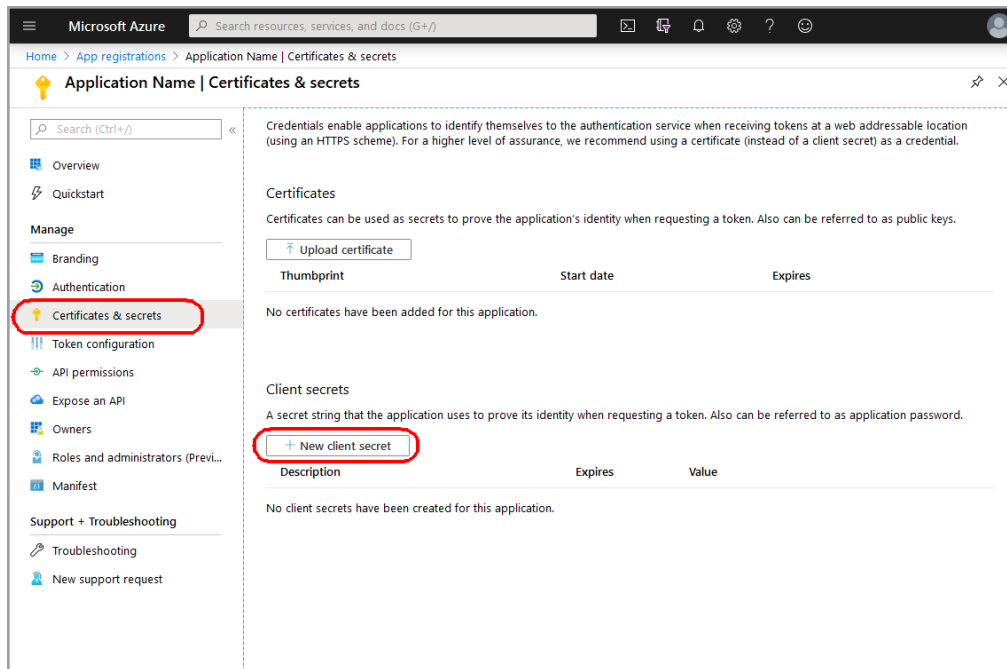
The screenshot shows the 'Register an application' page in the Microsoft Azure portal. The page has a dark header with the 'Microsoft Azure' logo and a search bar. Below the header, the breadcrumb trail is 'Home > App registrations > Register an application'. The main content area is titled 'Register an application' and contains the following sections:

- Name:** A text input field labeled 'Application Name' with a green checkmark icon to its right.
- Supported account types:** A section titled 'Who can use this application or access this API?' with three radio button options:
  - ☐ Accounts in this organizational directory only (Single tenant)
  - ☒ Accounts in any organizational directory (Any Azure AD directory - Multitenant)
  - ☐ Accounts in any organizational directory (Any Azure AD directory - Multitenant) and personal Microsoft accounts (e.g. Skype, Xbox)A link 'Help me choose...' is located below the radio buttons.
- Redirect URI (optional):** A section titled 'We'll return the authentication response to this URI after successfully authenticating the user. Providing this now is optional and it can be changed later, but a value is required for most authentication scenarios.' It contains a dropdown menu set to 'Web' and a text input field containing 'https://mywebapp.companyname.com' with a green checkmark icon to its right.

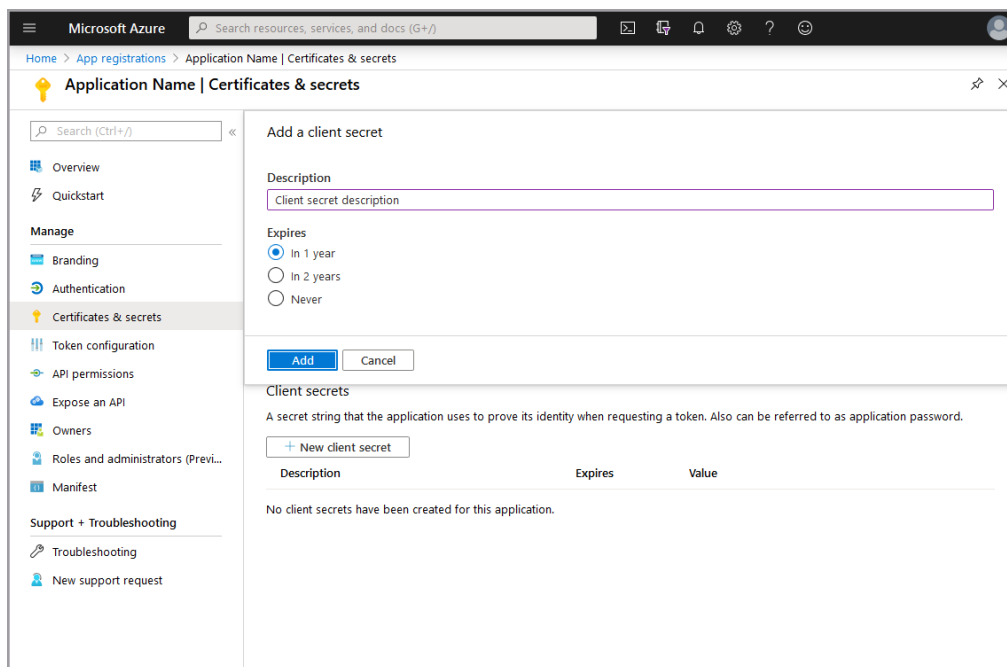
At the bottom of the form, there is a link 'By proceeding, you agree to the Microsoft Platform Policies' and a blue 'Register' button.

- **Name**—Enter a unique name for the application.
  - **Supported account types**—Select **Accounts in any organizational directory (Any Azure AD directory - Multitenant)**.
  - **Redirect URI**—If desired, you can select **Web** and enter a valid URL. The URL you enter will receive the authentication response after successful user authentication. You do not need to provide a redirect. If you do provide it, you can change it later. For more information, see <https://docs.microsoft.com/en-us/azure/active-directory/develop/reply-url>.
3. Click **Register**.

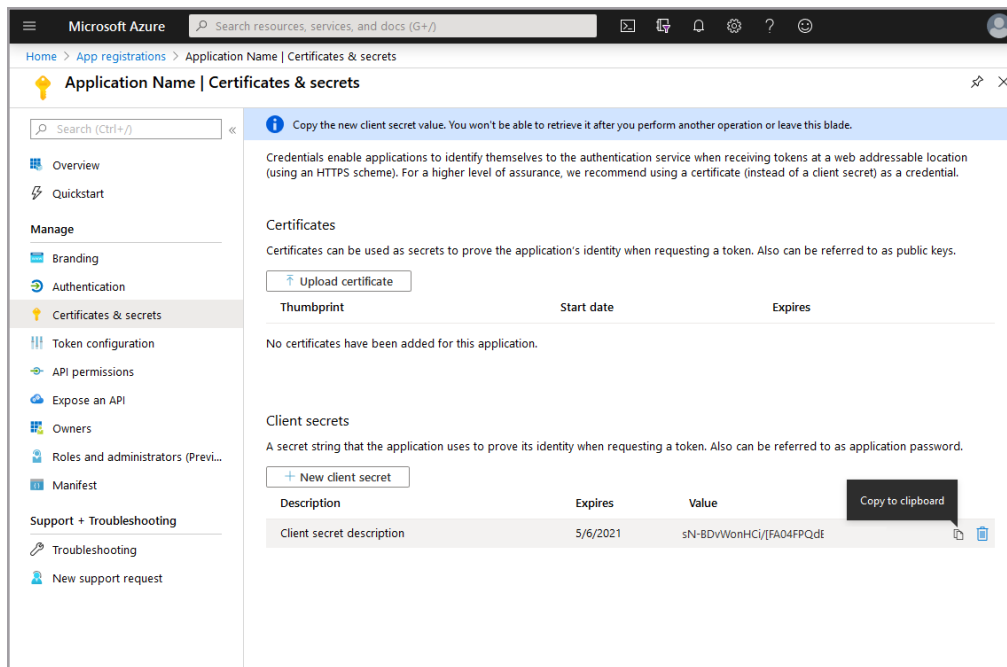
4. On the page for your new application, click **Certificates & secrets** in the left navigation menu and then click **New client secret**.



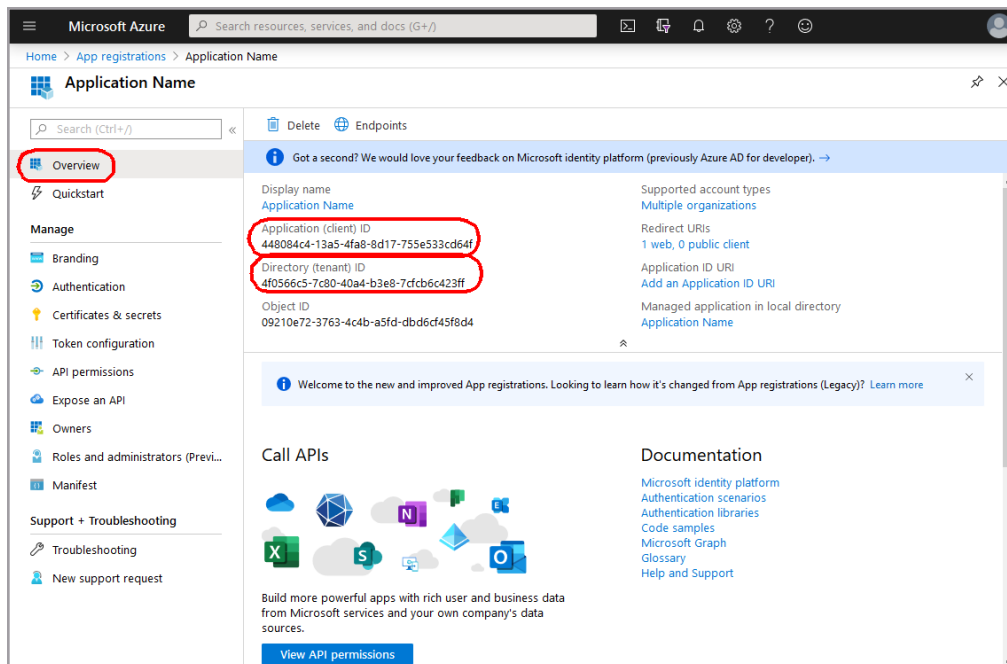
5. Specify the client secret description and an expiration, and then click **Add**.



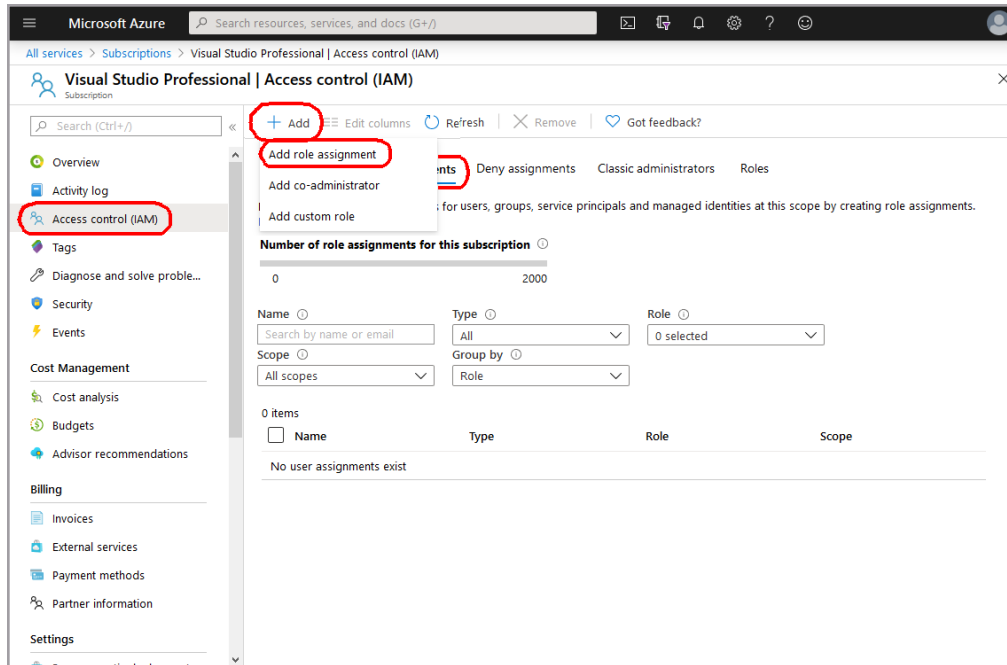
- In the table below **Client secrets**, click **Copy to clipboard** at the right of the table row for the client secret you just created. As the Azure portal mentions at the top, you must copy this value now. Once you leave this screen or do anything else, you will not have access to the client secret value. You will need this key for the **Client Secret** in Carbonite Migrate.



- In the left navigation menu, click **Overview**.
- Copy the **Application (client) ID** and the **Directory (tenant) ID**. They will be used for the **Client ID** and the **Tenant ID**, respectively, in Carbonite Migrate.



9. Assign the owner role to your application. Begin the process using one of the following methods.
  - Search for subscriptions in the top search bar.
  - Select **All services** from the fly out menu and then under **Categories**, select **All** and then **Subscriptions**.
10. From the list of subscriptions, click the subscription you will be using for Carbonite Migrate.
11. In the left navigation menu, click **Access control (IAM)**.
12. Select the **Role assignments** tab.
13. Click **Add, Add role assignment**.



14. Specify the role assignment values.

**Add role assignment** ✕


Role ⓘ  
Owner ⓘ

Assign access to ⓘ  
Azure AD user, group, or service prin... ▼

Select ⓘ  
Application Name

No users, groups, or service principals fo...

Selected members:

 Application Name [Remove](#)

Save Discard

- **Role**—Select **Owner**.
- **Assign access to**—Select **Azure AD user, group, or service principal**.
- **Select**—Enter the name of the application. This is the name you specified in step 2.

15. Click **Save**.

## Target requirements for Microsoft Azure Resource Manager (legacy)

In addition to the general *Target requirements* on page 15, a Microsoft Azure Resource Manager (legacy) environment must meet the following requirements.

- **Configuration**—You must register Carbonite Migrate as a consumer of Active Directory and assign it an owner role. During the registration, you will need to obtain a **Client ID**, **Client Secret**, and **Tenant ID** that will be required when creating an Azure Resource Manager (legacy) environment in Carbonite Migrate. See *Registering and assigning owner role in Microsoft Azure Resource Manager* on page 19 for details.
- **Networking**—A virtual network and subnet must be associated to the region you are migrating to.
- **Subscription**—You must have your Azure Active Directory subscription ID.
- **Templates**—During the migration setup, you must select an operating system template for each source operating system you are migrating. This template will be used to create the target server in the cloud.



## Target requirements for Microsoft Azure Classic

In addition to the general *Target requirements* on page 15, a Microsoft Azure Classic environment must meet the following requirements.

- **Subscription**—You must download and import an Azure subscription file.
- **Templates**—During the migration setup, you must select an operating system template for each source operating system you are migrating. This template will be used to create the target server in the cloud.



Each time you use a Windows template, the size of the system disk created on the target may be different than the size of the system disk on the template based on a calculation. The calculation is the system disk space in use on the source plus minimum amounts of additional free space required by Carbonite Migrate based on your operating system (10 GB for 2008 R2, 14 GB for 2012, and 15 GB for 2012 R2).

- If the calculation is smaller than the template system disk size, the template disk size will be used on the target.
  - If the calculation is larger than the template system disk size, the calculation size will be used on the target.
  - If the calculation is equal to or larger than the allowed maximum Azure size, the migration will not be allowed.
-

## Target requirements for custom environments

A custom environment as a target environment is only supported when the target server already exists. Carbonite Migrate will not create any virtual machines or disks, configure any target access points, open firewall ports, or other automated tasks that are usually performed when Carbonite Migrate creates the target for you. Therefore, the target must meet the following requirements, in addition to the general *Target requirements* on page 15.



Due to limitations provisioning storage for Linux servers in the cloud, Linux source servers can only be migrated to custom target environments at this time. This means you will need to have an existing target server already provisioned and running in your cloud to use as a custom target.

---

### • Matching volumes

- **Windows**— For each volume you are migrating from the source, the target must have a matching volume. For example, if you are migrating volumes C: and D: on the source, the target cannot have volumes D: and E:. In this case, the target must also have volumes C: and D:. Additional target volumes are preserved and available after cutover with all data still accessible.
- **Linux**— There are no limits to the number of volumes you can migrate on the source, although you are bound by operating system limits.

For each non-system volume you are protecting on the source, the target must have a matching volume. For example, if you are migrating /data and /home on the source, the target must also have /data and /home. Additional target volumes are preserved and available after cutover with all data still accessible.

The system volumes / and /boot do not have this matching volume limitation. If you have / and /boot on different volumes on the source, they can exist on a single volume on the target. If you have / and /boot on the same volume on the source, they can exist on different volumes on the target.

### • Operating system

- **Windows**— The source and target must have the same operating system. For example, you cannot have Windows 2008 on the source and Windows 2012 on the target. The two servers do not have to have the same level of service pack or hotfix. Windows 2008 and 2012 and their R2 releases are considered different operating systems. Therefore, you cannot have Windows 2008 or 2012 on the source and the corresponding R2 version on the target. The Windows edition (Standard, Enterprise, and so on) does not have to be the same.
- **Linux**— The source and target must have the same distribution and major version. For example, you cannot have a Red Hat version 5.11 source failing over to a Red Hat version 6.10 target. The two servers do not have to have the same minor version. For example, you can failover Red Hat version 6.9 to Red Hat version 6.10.
- **Disk space**— The target must have enough space to store the data from the source. If you do not have enough disk space, the migration will fail. This amount of disk space will depend on the applications and data files you are migrating. The more data you are migrating, the more disk space you will need. The target must also have enough space to store, process, and apply the

source's system state data. The size of the system state will depend on the operating system and architecture.

- **Windows**—A copy of the source system state data will be staged on the target boot volume in a folder called Staging-SSM. You can predict (approximately) how much space you will need in this staging folder by calculating the size of the following folders on your boot volume.
  - Documents and Settings
  - Program Files
  - Program Files (x86)
  - Program Data
  - Windows
  - Users

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

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

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

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

- **Linux**—A copy of the source data and system state will be staged on the target in a /dtstaging location for each mount point. For example, / will be staged in /dtstaging and /boot will be staged in /boot/dtstaging. You can predict how much space you will need in the staging folders by the amount of used space on the source.
- **Other target compatibility requirements**—There are other requirements for a compatible existing target that are detailed in the Carbonite Migrate for Windows and Linux documentation. See the [Documentation Library](#) and select the Carbonite Migrate User's Guide that matches your operating system. Inside the *User's Guide*, look under *Full server migration*, and then in *Full server migration requirements for a Target compatibility* section. This section details how an existing target server must be compatible with the source you are migrating.
- **Remote access**—The existing target must allow remote access. For Windows servers, the target must allow WinRM connections. For Linux servers, the target must allow SSH connections.
- **Connectivity**—The existing target must be accessible using the IP address you will provide for the migration. It also must be able to access the source using any of source's known IP addresses. The source machine must be able to access the target using the same IP you will provide for the migration.
- **Firewall**—For an existing Windows target, ports 6320, 6325, and 6326 must be open for UDP and TCP for both inbound and outbound traffic. For an existing Linux target, you must have ports 22, 1500, 1505, 1506, 6325, and 6326 open for both UDP and TCP for both inbound and outbound traffic, and you must also have port 1501 open for localhost communication.

- **Carbonite Migrate installation**—Because Carbonite Migrate is not automating creation of the target server in a custom environment, you must have Carbonite Migrate already installed on the existing target before you begin the migration. The target server must have the same packages and services installed as the source servers you are migrating. Carbonite Migrate will check for and, if necessary, install the required packages and services needed to install and use Carbonite Migrate. This process requires Internet access.

# Proxy requirements

If your source or target does not have a public IP address, you will have to use a proxy server for communication between the source or target and the Carbonite Migrate service. The proxy server is not used for communication between the source and target.

The proxy server must meet the following requirements.

- **Operating system**—The proxy server can be run on any of the following operating systems.
  - Windows 2019 Standard or Datacenter
  - Windows 2016 Standard or Datacenter



Server Core is not supported for any of these operating systems.

---

- **System memory**—The minimum system memory on the proxy server is 1 GB.
- **Disk space for program files**—The proxy server requires 150 MB of disk space for the program files.
- **Server where proxy is installed**—Do not install the proxy server on your source server.
- **Location where proxy is installed**— The proxy server cannot be located behind an Internet proxy. The proxy server must be directly connected to the Internet.
- **Microsoft .NET Framework**—Microsoft .NET Framework version 4.6.2 or later is required on the proxy server.



If an existing proxy server is not running Microsoft .NET 4.6.2 or later, you will need to manually upgrade to Microsoft .NET 4.6.2 or later before your proxy server will upgrade to the latest version of Carbonite Migrate. After the .NET upgrade, reboot the proxy server and it will automatically connect to Carbonite Migrate and upgrade.

---

# Carbonite Migrate interface requirements

Carbonite Migrate is the online service and interface that you use to execute and manage the migration. You must meet the following requirements to use the interface.

- **Web browser**—Any recent web browser can be used to access Carbonite Migrate, however the browser must have JavaScript enabled for the site to work correctly. If you are using Internet Explorer version 10 or later, make sure you are not running Compatibility Mode.
- **Carbonite Migrate to server communication**—If your source or target does not have a public IP address, you will have to use a proxy server for communication between the source or target and the Carbonite Migrate service. The proxy server is not used for communication between the source and target. See *Proxy requirements* on page 29 for details.

---

## Chapter 3 Getting started

Before you get started, make sure you have reviewed the *Requirements* on page 10 and have access to the *Carbonite Migrate interface* on page 32. Then complete the following tasks, in order.

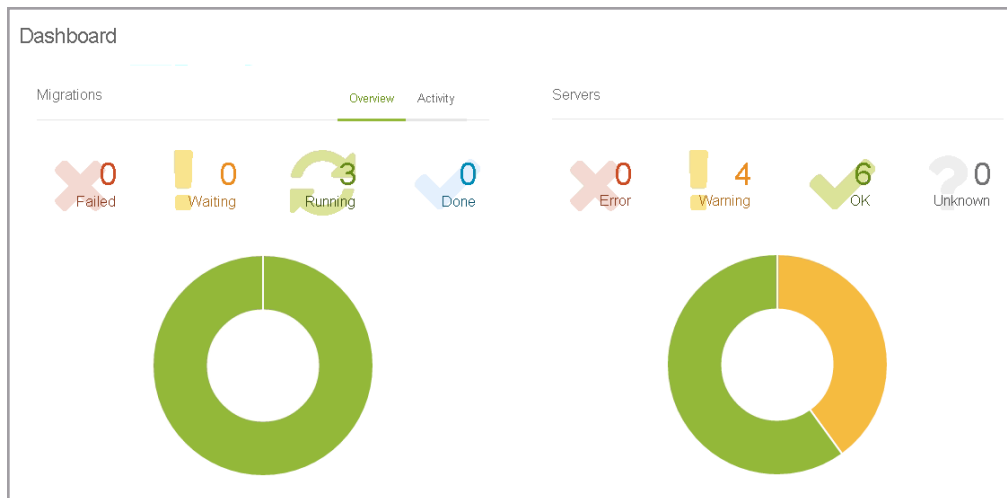
1. **Create environments**—Environments contain the information needed to connect to your servers. You will need at least two environments. One environment will be for the source so you can discover all of the servers you want to migrate. The second environment will be for the target so you can connect to your target cloud. If you are migrating from or to multiple environments, you will need an environment for each type. See *Environments* on page 36 for more details.
2. **Discover servers**—Once you have your source environment created, you can discover the source servers you want to migrate. See *Servers* on page 58 for more details.
3. **Migrate servers**—After your environments are created and your source servers discovered, you can migrate them. See *Migrations* on page 64 for more details.

---

## Chapter 4 Carbonite Migrate interface

You can access the Carbonite Migrate interface at <https://migrate.doubletake.com>.

- **First time access**—Use these steps if it is your first time accessing the service.
  - a. Click the **Create an account** link.
  - b. Provide the requested information, noting that the password must be at least six characters long and must contain characters from at least three of the following categories: uppercase letters, lowercase letters, numerals, or non-alphanumeric characters.
  - c. You must agree to the license agreement, terms of use, and privacy policy in order to use the service.
  - d. Click **Register**.
  - e. You will receive a verification email from [noreply@carbonite.com](mailto:noreply@carbonite.com). Click the **Verify Email** link in the message to complete your sign up.
- **Log In**—Use the email address and password you specified when first signing up for the service.
- **Dashboard**—Each time you access Carbonite Migrate, you will see a dashboard page, which is the highest level overview.



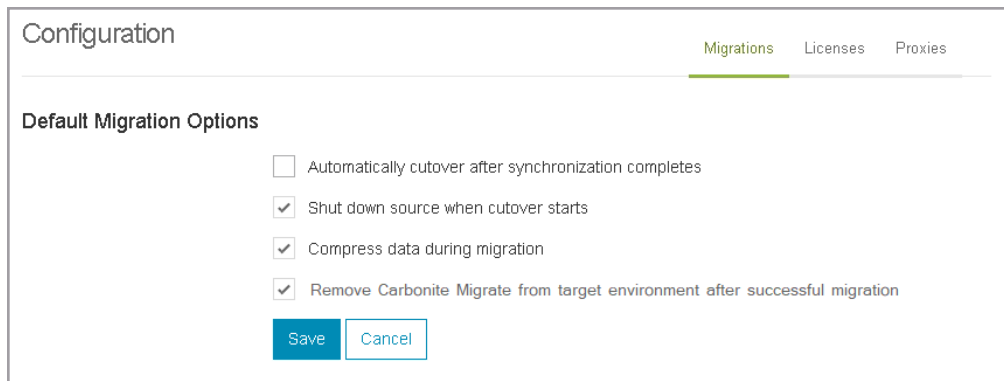
- **Migrations, Overview**—On the **Overview** tab of the **Migrations** section, you will see the number of migrations in each of the various states. Hover over colored areas of the large circle to see the number of migrations that are in that state. Click on a state name above the circle to go the **Migrations** page. Only those migrations that are in the state clicked will be displayed on the **Migrations** page. Clear the state name from the search box to see all migrations.
- **Migrations, Activity**—On the **Activity** tab of the **Migrations** section, you will see the migration activity log.
- **Servers**—On the Servers sections, you will see the number of servers in each of the various states. This is for servers in all of your environments. Hover over colored areas of the large circle to see the number of servers that are in that state. Click on a state name above the circle to go the **Servers** page. Only those servers that are in the state clicked will be displayed on the **Servers** page. Clear the state name from the search box to see all servers.



- **Alerts**—No matter which page you are on in the interface, you can access a list of alerts from the bell icon near the upper right corner of the page. The number of alerts will be displayed in a circle over the bell. If you have only information alerts, the circle will be green. If you have one or more warning alerts, the circle will be yellow. If you have one or more error alerts, the circle 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 alerts and the time they were generated. You can dismiss individual alerts or all of the alerts in the current list.
- **User profile**—Under your sign in name, in the upper right corner, is a menu of options.
  - **My Account**—Select this option to access your user profile. This is the account information you provided when you signed up as well as a few settings specific to your account.
    - **Delete**—The trash icon allows you to delete your Carbonite Migrate account. You will receive an email requiring you to confirm the deletion.
    - **Details**—This is your profile and contact information. You can also modify your account credentials here.
    - **Preferences**—You have several display options as well as email configuration.
      - **Show What's New on startup**—Specify if you want to show or hide the page that lists new features each time you access the Carbonite Migrate. You can view this page at any time by clicking **View**.
      - **Email Prefix**—By default, the subject line of email alerts sent to your account email will be prefaced with Carbonite Migrate Alert. This prefix allows you to recognize and filter emails specific to Carbonite Migrate. You can change or remove the prefix as desired.
      - **Email Alerts**—Select the type and level of alerts that you want to receive as email messages. If you do not select any type or level, you will not receive alerts as email messages. You will still get alerts in the Carbonite Migrate whether email alerts are enabled or disabled.
    - **Activity Log**—This is the log of activity that has occurred during your current user session. The log will be cleared when you sign out.
  - **Downloads**—Select this option to access download files.
    - **PowerShell SDK**—Carbonite Migrate includes Windows PowerShell cmdlets that you can use to control Carbonite Migrate. This menu option takes you to the Carbonite SDK site where you can download a zip file that contains example PowerShell scripts that you can adapt to your environment.
    - **Proxy Server**—If you need to install and register a proxy server, click **Download**. See *Proxy server* on page 57 for more details.
  - **Documentation**—Select this option to access the Carbonite Migrate *User's Guide* from the web.
  - **Log Out**—Select this option to manually sign out of Carbonite Migrate. If there is a period of inactivity, you will automatically be signed out. Any migrations that you have started will continue to run, even when you are signed out.

# Configuration

The **Configuration** page allows you to configuration your default migration options, manage your Carbonite Migrate licenses, and manage your proxy servers.



The screenshot shows the 'Configuration' page with the 'Migrations' tab selected. Under 'Default Migration Options', there are four checkboxes: 'Automatically cutover after synchronization completes' (unchecked), 'Shut down source when cutover starts' (checked), 'Compress data during migration' (checked), and 'Remove Carbonite Migrate from target environment after successful migration' (checked). At the bottom are 'Save' and 'Cancel' buttons.

- **Migrations**—You can set the default options that will be used. You can override these default settings when creating a migration.
  - **Default options**—These settings are the default migration options.
    - **Automatically cutover after synchronization completes**—When selected, cutover will begin automatically as soon as synchronization is complete. If this option is not selected, the migration process will continue replicating data changes but will not begin the cutover step until you click **Cutover** to continue the migration process.
    - **Shut down source when cutover starts**—When selected, the source will be shut down automatically when cutover begins. In most cases, you should shut down the source because conflicts can occur if you have two servers with the same name on the same network. Shutting down the source server allows the target to come up without conflicts. If this option is not selected, the source will remain online when cutover begins.
    - **Compress data during migration**—When selected, data will be compressed as it is sent across the wire from the source to the target. Compression can reduce the amount of bandwidth consumed as well as the time that it takes to synchronize all of the data to the target. If this option is not selected, the data is not compressed.
    - **Remove Carbonite Migrate from target environment after successful migration**—When selected, Carbonite Migrate will be uninstalled and any security groups needed for the migration, if created, are removed. This option is not available for custom target environments.



If Amazon Web Services cannot delete the security group because of dependent objects, you will need to manually disassociate or remove the security group from all dependent objects then manually delete it.

If you are using Amazon Web Services Administrator credentials and are migrating Windows 2008 R2 or Windows 2016 to Microsoft Azure (Resource Manager, Resource Manager (legacy), or Classic), Carbonite Migrate may

not be uninstalled during the cleanup process. In this case, you should not select the option to remove Carbonite Migrate after successful migration. To complete the uninstall, you need to remove the registry key HKEY\_CLASSES\_ROOT\Installer\Products\F64ABE42F8D664043BBEEE7E95CA9077\Transform from the target server after cutover, and then you can manually uninstall Carbonite Migrate.

---

- **Licenses**—Licenses will be listed with the version of the license, the used and available remaining quantities, and the date and time the license was updated. Click **Add License** to add another license to your environment. Click **Delete** to remove an existing license.
- **Proxies**—Proxy servers will be listed. Click **Delete** to remove an existing proxy server. If you need to install and register a proxy server, click **Download Installation File**. See *Proxy server* on page 57 for more details.

---

## Chapter 5 Environments

Environments contain the information needed to connect to your servers. You will need at least two environments. One environment will be for the source so you can discover all of the servers you want to migrate. The second environment will be for the target so you can connect to your target cloud. If you are migrating from or to multiple environments, you will need an environment for each type. The following environment types are supported. See *Target requirements* on page 15 for links to details on using each of these environments.

Environment Type	Supported Source Environments	Supported Target Environments
Amazon Web Services (AWS)	X	X
Google Cloud Platform	X	X
Microsoft Azure Resource Manager <sup>1</sup>		X
Microsoft Azure Resource Manager (legacy)		X
Microsoft Azure Classic		X
VMware vCloud Director	X	
VMware vSphere	X	
Active Directory	X	
Custom	X	X <sup>2</sup>

1. Use the Microsoft Azure Resource Manager target type if you have a migrate project under Azure Migrate in the Azure portal. A migrate project is used to store discovery, assessment, and migration metadata reported by the on-premise environment. If you do not have a migrate project, you can learn more about them at <https://azure.microsoft.com>. For migrations to Azure Resource Manager without using Azure Migrate projects, use the Microsoft Azure Resource Manager (legacy) target environment type.
2. A custom environment as a target environment is only supported when the target server already exists. See *Target requirements for custom environments* on page 26 for complete details on the requirements for using an existing server in a custom target environment.

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.

Environments

+ Add Proxy Delete

0 Selected Select All Clear All Selections 20 per page Search

<input type="checkbox"/>	Name ^	Environment Type ↕	Provider ↕	Proxy Server ↕	
<input type="checkbox"/>	Active Directory	Source	Active Directory	WIN2012R2	Actions ▾
<input type="checkbox"/>	Amazon Web Services	Source/Target	Amazon Web Services	WIN2012R2	Actions ▾
<input type="checkbox"/>	Custom	Source/Target	Custom	WIN2012R2	Actions ▾
<input type="checkbox"/>	Microsoft Azure Classic	Target	Microsoft Azure Classic	WIN2012R2	Actions ▾
<input type="checkbox"/>	Microsoft Azure Resource Manager	Target	Microsoft Azure Resource Manager	WIN2012R2	Actions ▾
<input type="checkbox"/>	OpenStack	Target	OpenStack	WIN2012R2	Actions ▾
<input type="checkbox"/>	VMware vCloud Director	Source/Target	VMware vCloud Director	WIN2012R2	Actions ▾
<input type="checkbox"/>	VMware vSphere	Source/Target	VMware vSphere	WIN2012R2	Actions ▾

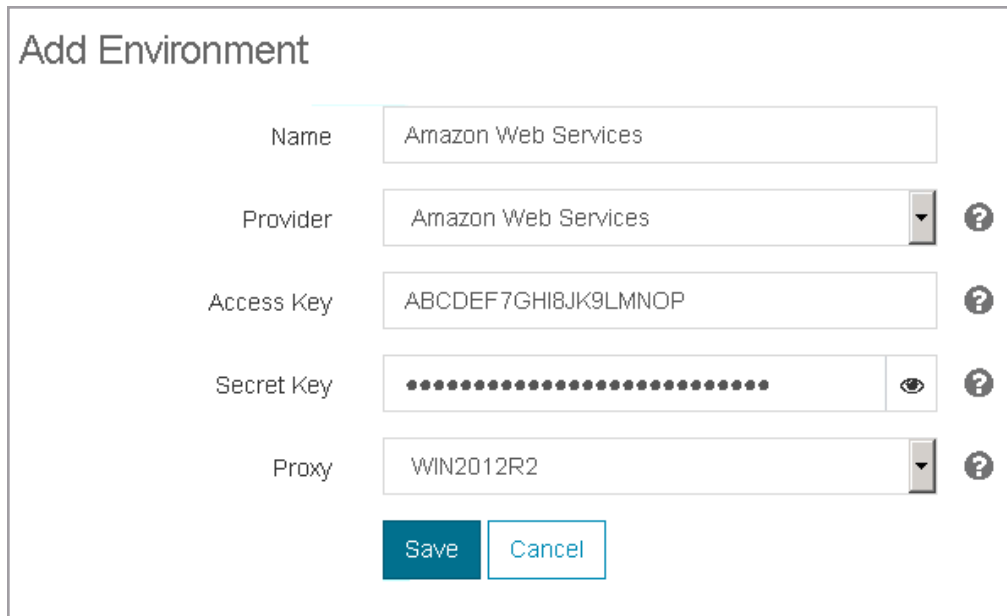
- **Toolbar and table controls**—The following controls are available on the toolbar and for the table on the **Environments** page.
  - **Add**—Click this button to add a new environment. See one of the following links for details on adding that environment type.
    - *Creating an Amazon Web Services environment* on page 39
    - *Creating a Google Cloud Platform environment* on page 43
    - *Creating a Microsoft Azure Resource Manager environment* on page 45
    - *Creating a Microsoft Azure Resource Manager (legacy) environment* on page 47
    - *Creating a Microsoft Azure Classic environment* on page 49
    - *Creating a VMware vCloud Director environment* on page 51
    - *Creating a VMware vSphere environment* on page 53
    - *Creating an Active Directory environment* on page 54
    - *Creating a custom environment* on page 56
  - **Proxy**—Select this option to change the proxy server assigned to the selected environments. See *Proxy server* on page 57 for more details on proxy servers.
  - **Delete**—Click this button to delete the selected environments. You cannot delete an environment that has a migration in progress.
  - **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.
  - **Search**—Text entered in the search 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 and discovery log for that environment.
  - **Settings**—You can modify the name of an environment, the information for the environment, or the proxy for the environment. For details on each of the environment settings, see the topic for creating that environment type. For details on proxy servers, see *Proxy server* on page 57.
  - **Discovery Log**—For source environments, this log lists the discoveries completed for this environment. For each log entry, you can click the right arrow to see more detailed information, if any, about that specific discovery process.
- **Single environment controls**—The following menu options are available in the **Actions** menu.
  - **Discovery Log**—For source environments, select this option to see a log of discoveries completed for this environment. For each log entry, you can click the right arrow to see more detailed information, if any, about that specific discovery process.
  - **Discover**—For source environments, select this option to discover servers in the environment. See *Discovering servers in non-custom environments* on page 60 or *Discovering servers in custom environments* on page 61 depending on the type of environment you are using.
  - **Proxy**—Select this option to change the proxy server assigned to the environment. See *Proxy server* on page 57 for more details on proxy servers.
  - **Delete**—Select this option to delete the environment. You cannot delete an environment that has a migration in progress.

# Creating an Amazon Web Services environment

Amazon Web Services (AWS) can be used as either a source or target environment. Make sure you have reviewed the *Source requirements* on page 11 and *Target requirements for Amazon Web Services* on page 16.

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



- **Name**—Specify a unique name for this environment. The name is limited to 64 characters.
  - **Provider**—Select **Amazon Web Services**.
  - **Access key**—Specify your AWS access key. This key is in your AWS account information. (In **IAM, Users**, select a user, then select **Security Credentials**.) The minimum permissions required for the access key for Carbonite Migrate are different depending on if you want to use the environment as a source or target. See *Minimum permissions for Amazon Web Services* on page 41 for details.
  - **Secret key**—Specify your AWS secret key. If you do not have your secret key, you will have to create a new one due to AWS security. See AWS documentation for creating a new key.
  - **Proxy**—If you are using a proxy server, select it from the list, otherwise, select **No proxy**. If you need to use a proxy server but have not installed and registered it yet, see *Proxy server* on page 57 for more details.
3. When you have identified your environment, click **Save**.



Generally, you only need to create one environment for each **Provider** type. You can perform multiple migrations for that environment. If you have environments defined by unique criteria, then you need to create multiple environments for each **Provider** type that requires unique criteria.

Because Carbonite Migrate cannot determine when credentials have changed, you will receive a warning notification to review and update your environment credentials every 30 days.

---



## Minimum permissions for Amazon Web Services

The minimum permissions required for the access key for Carbonite Migrate are different depending on if you want to use the environment as a source or target.

- **Source permissions**—The following permissions are required for a source environment.
  - AuthorizeSecurityGroupEgress
  - AuthorizeSecurityGroupIngress
  - CreateSecurityGroup
  - DescribeInstanceStatus
  - DescribeInstances
  - DescribeRegions
  - DescribeSecurityGroups
  - ModifyInstanceAttribute

To attach a policy to a specific user or group, access AWS using your administrator account. Specify a policy like the one below. (To specify the policy, go to **AWS Identify and Access Management (IAM) service, User/Group, Permission, Inline Policies, Create User/Group policy, Custom Policy.**)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:AuthorizeSecurityGroupEgress",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:CreateSecurityGroup",
        "ec2:DescribeInstanceStatus",
        "ec2:DescribeInstances",
        "ec2:DescribeRegions",
        "ec2:DescribeSecurityGroups",
        "ec2:ModifyInstanceAttribute"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

- **Target permissions**—Full control is required for a target environment.

To attach a policy to a specific user or group, access AWS using your administrator account. Specify a policy like the one below. (To specify the policy, go to **AWS Identify and Access**

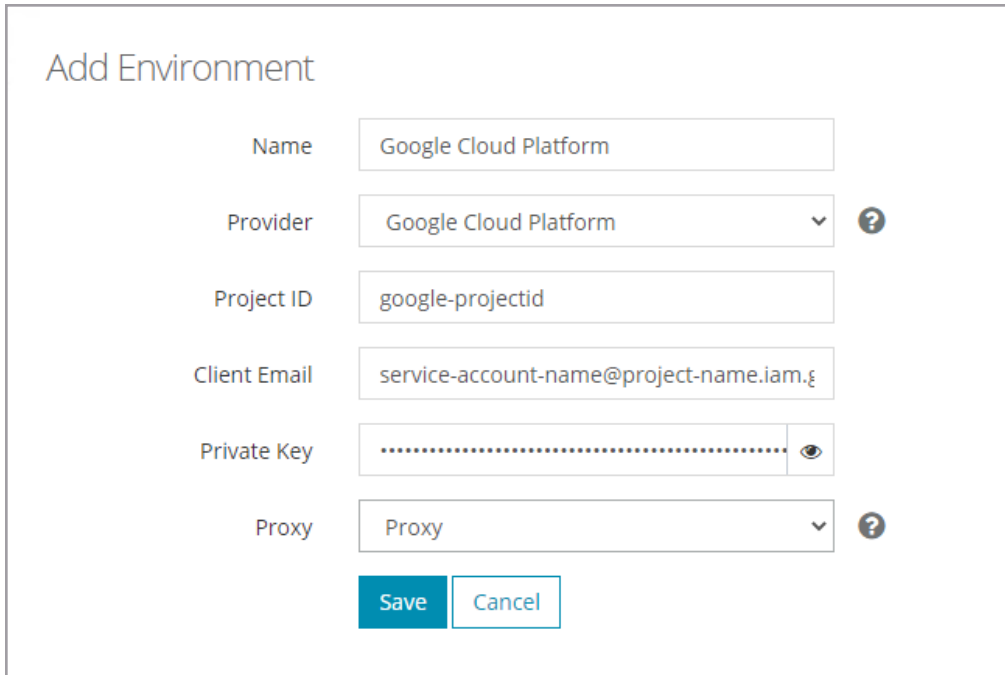
**Management (IAM) service, User/Group, Permission, Inline Policies, Create User/Group policy, Custom Policy.)**

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:*"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

# Creating a Google Cloud Platform environment

Google Cloud Platform (GCP) can be used as either a source or target environment. Make sure you have reviewed the *Source requirements* on page 11 and *Target requirements for Google Cloud Platform* on page 17.

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



- **Name**—Specify a unique name for this environment. The name is limited to 64 characters.
- **Provider**—Select **Google Cloud Platform**.
- **Project ID**—Specify your Google Project ID.
- **Client Email**—Specify the service account email associated with the specified Google Cloud project.
- **Private Key**—Specify the private key associated with the service account. The private key is everything between the quotation marks (") in the .json file that you downloaded from Google. The private key includes the comment and line return characters. For example, your private key may look similar to the following, but your key will be much longer (over 1500 characters).

```
"private_key": "-----BEGIN PRIVATE KEY-----  
\nMIIIEvQ0BAQjMjV+9aYG/nFDaHJqmkQ/kSQPXE2a7uw1abf\nAzf9KGb/1m+iz  
Xykta0j2qc=\n-----END PRIVATE KEY-----\n",
```

In this example you would use the following for the Private Key.

```
-----BEGIN PRIVATE KEY-----  
\nMIIIEvQ0BAQjMjV+9aYG/nFDaHJqmkQ/kSQPXE2a7uw1abf\nAzf9KGb/1m+iz  
Xykta0j2qc=\n-----END PRIVATE KEY-----\n
```

- **Proxy**—If you are using a proxy server, select it from the list, otherwise, select **No proxy**. If you need to use a proxy server but have not installed and registered it yet, see *Proxy server* on page 57 for more details.
3. When you have identified your environment, click **Save**.



Generally, you only need to create one environment for each **Provider** type. You can perform multiple migrations for that environment. If you have environments defined by unique criteria, then you need to create multiple environments for each **Provider** type that requires unique criteria.

---

# Creating a Microsoft Azure Resource Manager environment

Microsoft Azure Resource Manager can be used only as a target environment. Make sure you have reviewed the *Target requirements for Microsoft Azure Resource Manager* on page 18. Also make sure you have obtained a **Client ID**, **Client Secret**, and **Tenant ID**. See *Registering and assigning owner role in Microsoft Azure Resource Manager* on page 19

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

The screenshot shows a web form titled "Add Environment". It contains the following fields and values:

- Name:** Microsoft Azure Resource Manager
- Provider:** Microsoft Azure Resource Manager (dropdown menu)
- Client ID:** 12a3bc45-d67e-8901-abcd-23e4a567890b
- Client Secret:** [Masked with 12 dots]
- Tenant ID:** b09765ae-d54621-cb-0a98-ed76543210c9
- Subscription ID:** SubID
- Proxy:** Proxy (dropdown menu)

At the bottom of the form are two buttons: "Save" (in blue) and "Cancel" (in white with a blue border). Each input field has a small question mark icon to its right for help.

- **Name**—Specify a unique name for this environment. The name is limited to 64 characters.
- **Provider**—Select **Microsoft Azure Resource Manager**.
- **Client ID**—Specify the client ID obtained when registering Carbonite Migrate as a consumer in Active Directory. See the *Target requirements for Microsoft Azure Resource Manager* on page 18 for more details.
- **Client Secret**—Specify the key you added when registering Carbonite Migrate as a consumer in Active Directory. See the *Target requirements for Microsoft Azure Resource Manager* on page 18 for more details.
- **Tenant ID**—Specify the tenant ID obtained when registering Carbonite Migrate as a consumer in Active Directory. See the *Target requirements for Microsoft Azure Resource Manager* on page 18 for more details.
- **Subscription ID**—Specify your Azure Active Directory subscription ID.

- **Proxy**—If you are using a proxy server, select it from the list, otherwise, select **No proxy**. If you need to use a proxy server but have not installed and registered it yet, see *Proxy server* on page 57 for more details.

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

---



Generally, you only need to create one environment for each **Provider** type. You can perform multiple migrations for that environment. If you have environments defined by unique criteria, then you need to create multiple environments for each **Provider** type that requires unique criteria.

Because Carbonite Migrate cannot determine when credentials have changed, you will receive a warning notification to review and update your environment credentials every 30 days.

---

# Creating a Microsoft Azure Resource Manager (legacy) environment

Microsoft Azure Resource Manager (legacy) can be used only as a target environment. Make sure you have reviewed the *Target requirements for Microsoft Azure Resource Manager (legacy)* on page 24. Also make sure you have obtained a **Client ID**, **Client Secret**, and **Tenant ID**. See *Registering and assigning owner role in Microsoft Azure Resource Manager* on page 19

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

The screenshot shows the 'Add Environment' form with the following fields and values:

Field	Value
Name	Microsoft Azure Resource Manager (legacy)
Provider	Microsoft Azure Resource Manager (legacy)
Client ID	12a3bc45-d67e-8901-abcd-23e4a567890b
Client Secret	.....
Tenant ID	b09765ae-d54621-cb-0a98-ed76543210c9
Subscription ID	SubID
Proxy	Proxy

Buttons: Save, Cancel

- **Name**—Specify a unique name for this environment. The name is limited to 64 characters.
- **Provider**—Select **Microsoft Azure Resource Manager (legacy)**.
- **Client ID**—Specify the client ID obtained when registering Carbonite Migrate as a consumer in Active Directory. See the *Target requirements for Microsoft Azure Resource Manager (legacy)* on page 24 for more details.
- **Client Secret**—Specify the key you added when registering Carbonite Migrate as a consumer in Active Directory. See the *Target requirements for Microsoft Azure Resource Manager (legacy)* on page 24 for more details.
- **Tenant ID**—Specify the tenant ID obtained when registering Carbonite Migrate as a consumer in Active Directory. See the *Target requirements for Microsoft Azure Resource Manager (legacy)* on page 24 for more details.
- **Subscription ID**—Specify your Azure Active Directory subscription ID.

- **Proxy**—If you are using a proxy server, select it from the list, otherwise, select **No proxy**. If you need to use a proxy server but have not installed and registered it yet, see *Proxy server* on page 57 for more details.

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

---



Generally, you only need to create one environment for each **Provider** type. You can perform multiple migrations for that environment. If you have environments defined by unique criteria, then you need to create multiple environments for each **Provider** type that requires unique criteria.

Because Carbonite Migrate cannot determine when credentials have changed, you will receive a warning notification to review and update your environment credentials every 30 days.

---



# Creating a Microsoft Azure Classic environment

Microsoft Azure Classic can be used only as a target environment. Make sure you have reviewed the *Target requirements for Microsoft Azure Classic* on page 25.

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

**Add Environment**

Name: Microsoft Azure Classic

Provider: Microsoft Azure Classic

Publish Settings:

1. [Download](#) subscription file from Azure account.
2. Import the subscription file you downloaded. The file will have a .publishsettings extension:

[Upload file](#) Visual Studio Premium.publishsettings

Subscription: CMC

Proxy: WIN2012R2

[Save](#) [Cancel](#)

- **Name**—Specify a unique name for this environment. The name is limited to 64 characters.
  - **Provider**—Select **Microsoft Azure Classic**.
  - **Publish Settings**—Import a subscription file from Microsoft Azure using the following steps.
    1. If you do not already have an Azure publish settings file that you can import, click **Download** to download one from Microsoft Azure, entering your Azure credentials if prompted.
    2. Click **Upload File** to locate and import the .publishsettings file.
    3. If you have multiple subscriptions in your Azure account, select the **Subscription** you want to use with Carbonite Migrate.
  - **Proxy**—If you are using a proxy server, select it from the list, otherwise, select **No proxy**. If you need to use a proxy server but have not installed and registered it yet, see *Proxy server* on page 57 for more details.
3. When you have identified your environment, click **Save**.



Generally, you only need to create one environment for each **Provider** type. You can perform multiple migrations for that environment. If you have environments defined by



unique criteria, then you need to create multiple environments for each **Provider** type that requires unique criteria.

Because Carbonite Migrate cannot determine when credentials have changed, you will receive a warning notification to review and update your environment credentials every 30 days.








---

# Creating a VMware vCloud Director environment

VMware vCloud Director can be used only as a source environment. Make sure you have reviewed the *Source requirements* on page 11.

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

### Add Environment

Name	<input type="text" value="VMware vCloud Director"/>	
Provider	<input type="text" value="VMware vCloud Director"/>	 
Endpoint/IP	<input type="text" value="112.47.12.7"/>	
Organization	<input type="text" value="OrgName"/>	
Username	<input type="text" value="administrator"/>	
Password	<input type="password" value="••••••••••"/>	
Proxy	<input type="text" value="WIN2012R2"/>	 
<div><input type="button" value="Save"/> <input type="button" value="Cancel"/></div>		

- **Name**—Specify a unique name for this environment. The name is limited to 64 characters.
  - **Provider**—Select **vCloud**.
  - **Endpoint/IP**—Specify the IP address or hostname of your vCloud environment.
  - **Organization**—Specify the organization name where your source servers are located.
  - **Username**—Specify a user to establish a session with vCloud Director.
  - **Password**—Specify the password associated with the user you have entered.
  - **Proxy**—If you are using a proxy server, select it from the list, otherwise, select **No proxy**. If you need to use a proxy server but have not installed and registered it yet, see *Proxy server* on page 57 for more details.
3. When you have identified your environment, click **Save**.



Generally, you only need to create one environment for each **Provider** type. You can perform multiple migrations for that environment. If you have environments defined by unique criteria, then you need to create multiple environments for each **Provider** type that requires unique criteria.

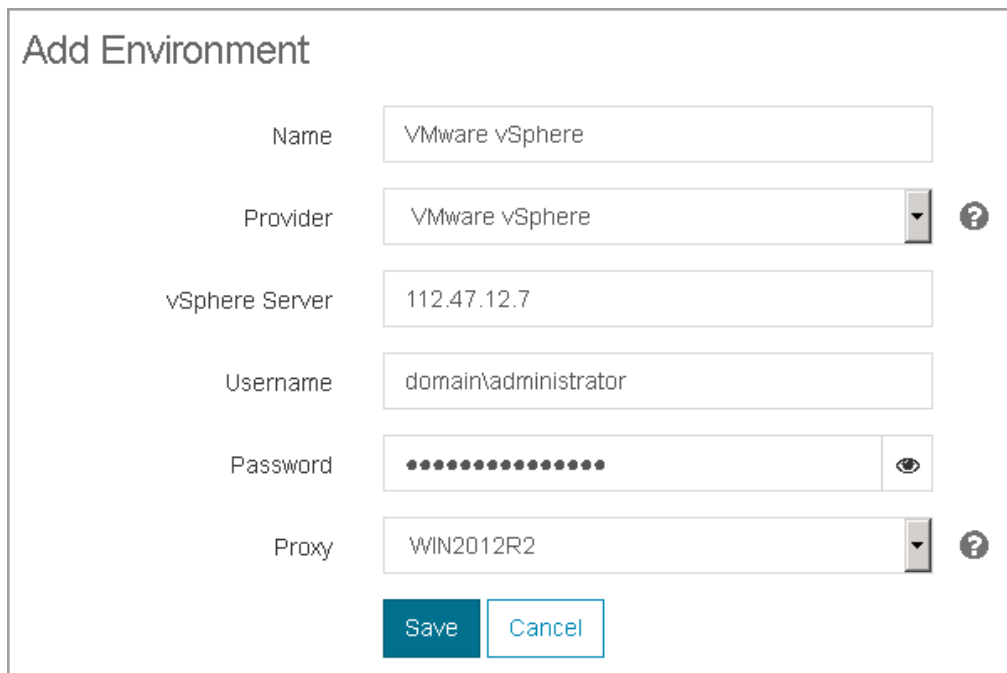
Because Carbonite Migrate cannot determine when credentials have changed, you will receive a warning notification to review and update your environment credentials every 30 days.

---

# Creating a VMware vSphere environment

VMware vSphere can be used only as a source environment. Make sure you have reviewed the *Source requirements* on page 11.

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



- **Name**—Specify a unique name for this environment. The name is limited to 64 characters.
  - **Provider**—Select **VMware vSphere**.
  - **vSphere Server**—Specify the IP address or host name of your vSphere or ESX server.
  - **Username**—Specify a user name for the vSphere or ESX server.
  - **Password**—Specify the password associated with the user you have entered.
  - **Proxy**—If you are using a proxy server, select it from the list, otherwise, select **No proxy**. If you need to use a proxy server but have not installed and registered it yet, see *Proxy server* on page 57 for more details.
3. When you have identified your environment, click **Save**.



Generally, you only need to create one environment for each **Provider** type. You can perform multiple migrations for that environment. If you have environments defined by unique criteria, then you need to create multiple environments for each **Provider** type that requires unique criteria.

Because Carbonite Migrate cannot determine when credentials have changed, you will receive a warning notification to review and update your environment credentials every 30 days.

# Creating an Active Directory environment

Active Directory can be used only as a source environment. Make sure you have reviewed the *Source requirements* on page 11.

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

**Add Environment**

Name: Active Directory

Provider: Active Directory ?

Domain Name: domain.com

LDAP Path: OU=SERVERS ?

Username: administrator

Password: [masked] [visibility icon]

Proxy: WIN2012R2 ?

Save Cancel

- **Name**—Specify a unique name for this environment. The name is limited to 64 characters.
- **Provider**—Select **Active Directory**.
- **Domain Name**—Specify the domain name.
- **LDAP Path**—If desired, specify the LDAP path to query a specific Active Directory object. Only the distinguished name of the object is queried. For example, if you want to query an entire domain name, domain.com, you do not need to specify an LDAP path. However, if you want to query Organization Unit (OU) SERVERS in domain.com, you can specify one of the following LDAP paths.

OU=SERVERS

OU=SERVERS,DC=DOMAIN,DC=COM

LDAP AD Path formats such as LDAP://dc01 or

LDAP://DC01/OU=SERVERS,DC=DOMAIN,DC=COM are also supported.

- **Username**—Specify a user name to establish a session with the Active Directory server.
- **Password**—Specify the password associated with the user you have entered.

- **Proxy**—If you are using a proxy server, select it from the list, otherwise, select **No proxy**. If you need to use a proxy server but have not installed and registered it yet, see *Proxy server* on page 57 for more details.

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

---



Generally, you only need to create one environment for each **Provider** type. You can perform multiple migrations for that environment. If you have environments defined by unique criteria, then you need to create multiple environments for each **Provider** type that requires unique criteria.

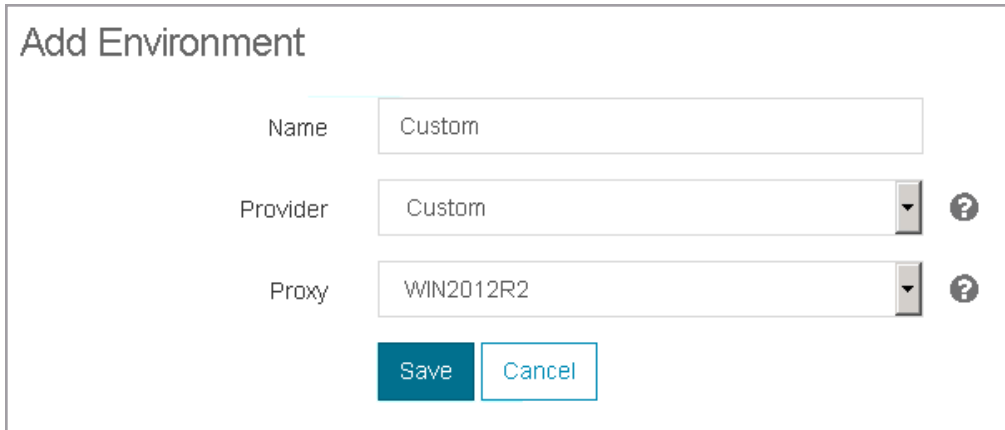
Because Carbonite Migrate cannot determine when credentials have changed, you will receive a warning notification to review and update your environment credentials every 30 days.

---

## Creating a custom environment

Custom environments can be used as either a source or target environment. Make sure you have reviewed the *Source requirements* on page 11 and *Target requirements for custom environments* on page 26.

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



The screenshot shows a dialog box titled "Add Environment". It has three input fields: "Name" with the text "Custom", "Provider" with a dropdown menu showing "Custom" and a help icon, and "Proxy" with a dropdown menu showing "WIN2012R2" and a help icon. At the bottom are "Save" and "Cancel" buttons.

- **Name**—Specify a unique name for this environment. The name is limited to 64 characters.
  - **Provider**—Select **Custom**.
  - **Proxy**—If you are using a proxy server, select it from the list, otherwise, select **No proxy**. If you need to use a proxy server but have not installed and registered it yet, see *Proxy server* on page 57 for more details.
3. When you have identified your environment, click **Save**.



Generally, you only need to create one environment for each **Provider** type. You can perform multiple migrations for that environment. If you have environments defined by unique criteria, then you need to create multiple environments for each **Provider** type that requires unique criteria.



# Proxy server

If your source or target does not have a public IP address, you will have to use a proxy server for communication between the source or target and the Carbonite Migrate service. The proxy server is not used for communication between the source and target. See *Proxy requirements* on page 29 for a complete list of the proxy server requirements.

Use the following instructions to configure a proxy server.

1. You can access the proxy server download from multiple locations.
  - Select **Downloads** from the menu under your sign in name in the upper right corner. Then click **Download** to the right of **Proxy Server**.
  - From the **Configuration** page, **Proxies** section, click **Download Installation File**.
2. Save and run the file.
3. Click **Install** to install the utility used to install and register the Carbonite Migrate proxy server software.
4. At the welcome page, click **Next**.
5. Review the license agreement. You must agree to the Terms of Service to continue the installation.
6. Click **Next** to continue.
7. If desired, change the location of the proxy server installation, and click **Next** to continue.
8. Click **Install** to begin the installation.
9. Enable the option to **Launch Proxy registration tool**. If you do not register the proxy now, you will need to run **Start, Apps, Carbonite Migrate Proxy, Carbonite Migrate Proxy Registration** before you can use the server as a Carbonite Migrate proxy.
10. **Close** the proxy server installation window.
11. On the proxy registration window, specify the server name that you want to register as the proxy server.
12. Specify your Carbonite Migrate account credentials. This will register your proxy server to your account.
13. Click **Connect**.
14. Once the proxy server is connected, you will see info on the **Information** tab and log entries on the **Logs** tab. You can see the information and logs at any time by re-opening the Carbonite Migrate Proxy Registration.
15. Click **Close** to close the proxy registration window.



The proxy server uses default service bus connectivity mode which means it auto-detects accessing the Carbonite Migrate service. First it will try using TCP on port 9354. If that is blocked, it will use REST over HTTPS on port 443. If that fails, it will use REST over HTTP on port 80. When accessing the Carbonite Migrate installation files to push to your servers, it will use HTTPS on port 443.

The proxy server cannot be located behind an Internet proxy. The proxy server must be directly connected to the Internet.

---

## Chapter 6 Servers

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.

<input type="checkbox"/>	Name ^	Status ^	Environment ^	IP Address ^	Platform ^	Location ^	Service Group ^	
<input type="checkbox"/>	ALPHA	Ready	VMware vSphere	172.31.206.200	Microsoft Windows Server 2008 R2 Enterprise	On Site	-	Actions v
<input type="checkbox"/>	BETA	Ready	VMware vSphere	172.31.206.201	Microsoft Windows Server 2008 R2 Enterprise	On Site	-	Actions v
<input type="checkbox"/>	DELTA	Ready	Custom	172.31.186.63	Microsoft Windows Server 2012 R2 Datacenter	Custom	-	Actions v
<input type="checkbox"/>	GAMMA	Ready	VMware vSphere	172.31.187.189	CentOS	On Site	-	Actions v

- **Toolbar and table controls**—The following controls are available on the toolbar and for the table on the **Servers** page.
  - **Discover**—Click this button to discover servers in one of your source environments. See *Discovering servers in non-custom environments* on page 60 or *Discovering servers in custom environments* on page 61 for details. Once a server is inventoried successfully, you can discover within the same environment again using different credentials, without the discovery process resetting the credentials of the successfully inventoried servers.
  - **Migrate**—Click this button to begin migration of the selected servers. See *Migrations* on page 64 for details.
  - **Credentials**—Click this button to specify credentials for the selected servers.
  - **Inventory**—Click this button to inventory the selected servers. Inventory gathers detailed information about the servers that could not be obtained from the initial discovery process, such as volumes, operating system, CPU, memory, and networking. You can click this button on a server that has already been inventoried to refresh and collect updated server information, such as changed IP addresses.
  - **Delete**—Click this button to delete the selected servers. You cannot delete a server that has a migration in progress.
  - **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.
  - **Search**—Text entered in the search 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 server to see the details for that server. See *Viewing server details* on page 63
- **Single server controls**—The following menu options are available by clicking the name of a server.
  - **Connect**—For a Windows server, select this option to open a Remote Desktop Connection to the server.
  - **Credentials**—Select this option to update the credentials used to access the server.
  - **Inventory**—Select this option to inventory the servers. Inventory gathers detailed information about the server that could not be obtained from the initial discovery process, such as volumes, operating system, CPU, memory, and networking. You can select this option on a server that has already been inventoried to refresh and collect updated server information, such as changed IP addresses.
  - **Migrate**—Select this option to migrate the server. Only servers in a ready status can be migrated. See *Migrations* on page 64 for more details.
  - **Delete**—Select this option to delete the server. You cannot delete a server that has a migration in progress.
- **Server status**—The **Status** column has different colors and icons to indicate the health of your servers.
  - **Green checkmark**—A green circle with a white checkmark indicates a good status.
  - **Yellow exclamation point**—A yellow circle with a white exclamation point indicates a pending or warning status.
  - **Red X**—A red circle with a white X indicates an error status.
  - **Black question mark**—A black circle with a white question mark indicates the status is unknown.

## Discovering servers in non-custom environments

1. Make sure you have created a source environment. See *Environments* on page 36 for details.
2. On the **Servers** page, click **Discover**, or on the **Environments** page, click **Discover** from the environment's **Actions** menu.
3. Select your **Environment**, if it is not already selected.

### Discover Servers

Enter the credentials used to sign in to servers in the selected environment. Servers for which these credentials are not valid will need to be inventoried separately.

Environment	VMware vCloud Director	▼
Credentials	Specify credentials...	▼
Username	administrator	
Password	••••••••••	👁

DiscoverCancel

4. Specify or select **Credentials** for the servers you are discovering. If the credentials you selected or specified are not valid for all servers, the servers will still be discovered but will be displayed in a credentials required state. You can update the credentials later from the **Servers** page.
5. Click **Discover**.



Once a server is inventoried successfully, you can discover within the same environment again using different credentials, without the discovery process resetting the credentials of the successfully inventoried servers.

---

## Discovering servers in custom environments

1. Make sure you have created a custom source environment. See *Creating a custom environment* on page 56 for details.
2. On the **Servers** page, click **Discover**, or on the **Environments** page, click **Discover** from the environment's **Actions** menu.
3. Select your custom **Environment**, if it is not already selected.
4. If you want to discover a single server by IP address or name, select **Server** for **Discover Type** and specify the following options.

### Discover Servers

Enter the credentials used to sign in to servers in the selected environment. Servers for which these credentials are not valid will need to be inventoried separately.

Environment	Custom
Discovery Type	Server
IP Address	172.31.206.200
Username	administrator
Password	••••••••
OS Type	Windows

DiscoverCancel

- **IP Address**—Specify the server's IP address or name. If you specify a computer name, Carbonite Migrate or the proxy server must be able to resolve it to an IP address.
  - **Username**—Specify a user for the server you are discovering. If the credentials are not valid, the server will still be discovered but will be displayed in a credentials required state. You can update the credentials later from the **Servers** page.
  - **Password**—Specify the password associated with the user you have entered.
  - **OS Type**—Specify the **Windows** or **Linux** operating system that is running on the server you are discovering.
5. If you want to discover multiple servers by a range of IP addresses, select **Network** for **Discover Type** and specify the following options.

## Discover Servers

Enter the credentials used to sign in to servers in the selected environment. Servers for which these credentials are not valid will need to be inventoried separately.

Environment	Custom
Discovery Type	Network
IP Network	172.31.206.200-172.31.206.230
Credentials	Specify credentials...
Username	administrator
Password	••••••••••••••••

Discover Cancel

- **IP Network**—Specify an IP address range or the network and subnet mask using one of these formats.
  - **Address range**—172.31.206.200 - 172.31.206.230
  - **Network and subnet mask**—172.31.206.200/24



Do not use a single IP address or range.

The maximum address range is 256 and the maximum subnet mask value is 24.

Only private IP addresses are supported.

If you need to discover multiple networks or ranges, repeat the discovery process.

- **Credentials**—Specify or select credentials for the servers you are discovering. If the credentials you selected or specified are not valid for all servers, the servers will still be discovered but will be displayed in a credentials required state. You can update the credentials later from the **Servers** page.

6. Click **Discover**.

## Viewing server details

On the **Servers** page, click on the name of a server. You will see three sections of server information.

Server Details		Connectivity	Properties	Disks
Name	ALPHA			
Username	administrator			
Public IP	172.31.206.200			
Public DNS	-			
Private IP	-			
Private DNS	-			

- **Connectivity**—This section contains information related to the server's connectivity to Carbonite Migrate.
- **Properties**—This section contains the server's properties.
- **Disks**—This section contains the server's disk information.

---

## Chapter 7 Migrations

On the **Migrations** page, you will find high-level information and controls for your migrations. You can sort the migration list by clicking on any column heading. Group migrations can be sorted by the group name or the individual migrations.

You can select multiple migrations on this page by selecting the checkbox to the left of a table row. A checkmark indicates the migration is selected.

<a href="#">+ Create</a> <a href="#">Delete</a> <a href="#">History</a> <a href="#">Cutover</a>		
0 Selected <a href="#">Select All</a> <a href="#">Clear All Selections</a> 20 per page <input type="text" value="Search"/>		
Group Name ^	Target Environment ^	Status ^
<input type="checkbox"/> Alpha and Beta migration	VMware vCloud Director	Warning <a href="#">Actions</a>
<input type="checkbox"/> ALPHA	ALPHA	Waiting to cutover <a href="#">Actions</a>
<input type="checkbox"/> BETA	BETA	Synchronizing - 89% <a href="#">Actions</a>
<input type="checkbox"/> Gamma migration	VMware vCloud Director	Ok <a href="#">Actions</a>
<input type="checkbox"/> GAMMA	GAMMA	Provisioning <a href="#">Actions</a>

- **Toolbar and table controls**—The following controls are available on the toolbar and for the table on the **Migrations** page.
  - **Create**—Click this button to create a new migration for the selected servers. See one of the following links for details on creating that migration type.
    - *Migrating to Amazon Web Services* on page 67
    - *Migrating to Google Cloud Platform* on page 73
    - *Migrating to Microsoft Azure Resource Manager* on page 79
    - *Migrating to Microsoft Azure Resource Manager (legacy)* on page 86
    - *Migrating to Microsoft Azure Classic* on page 92
    - *Migrating to a custom environment* on page 98



Due to limitations provisioning storage for Linux servers in the cloud, Linux source servers can only be migrated to custom target environments at this time. This means you will need to have an existing target server already provisioned and running in your cloud to use as a custom target.

- **Delete**—Click this button to delete the selected migrations. You will see an additional option to **Delete target resources created by the selected migrations**. Be careful when deleting the target resources. Only delete them if your migration failed or you do not want to use the target. Deleting the resources will delete everything that Carbonite Migrate created, including your target server.
- **History**—Click this button to see a history of your migrations. See *Viewing migration history* on page 103 for more details.



- **Cutover**—Click this button to cutover the selected migrations. All of the selected migrations must be in a **Waiting to cutover** state in order for this button to be available.
- **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.
- **Search**—Text entered in the search box will narrow the list displayed to only those rows that contain the search text.
- **Show migration groups**—Click this button to display migration group names. Individual migrations in a group are nested under their group name. You can sort at the group level by group, environment, or status. You can sort within a group by source, target, or status.
- **Hide migration groups**—Click this button to hide migration group names. Individual migrations will not be nested under their group name. Sorting is by source, target, or status.
- **Group migration controls**—If you are showing migration groups, you have menu options that are available for the migrations within that group.
  - **Cutover**—Select this option to cutover all of the migrations in the group. All of the selected migrations must be in a **Waiting to cutover** state in order for this option to be available.
  - **Delete**—Select this option to delete all of the migrations in the group. You will see an additional option to **Delete cloud resources created by selected migrations**. Be careful when deleting the cloud resources. Only delete them if your migration failed or you do not want to use the target. Deleting the cloud resources will delete everything in the target cloud, including your target server.
- **Hyperlink control**—Click the name of a migration to see details for that migration. See *Viewing migration details* on page 102.
- **Single migration controls**—The following menu options are available by clicking the name of a individual migration. (Expand any group migrations to see the individual migrations.)
  - **View log**—Select this option to see a log of the workflow steps for this migration.
  - **Connect Source**—For a Windows server, select this option to open a Remote Desktop Connection to the source server.
  - **Connect Target**—For a Windows server, select this option to open a Remote Desktop Connection to the target server.
  - **Resume**—This button is only displayed if a migration is in an error or paused state. Click this button to restart the migration from the point it failed or paused.
  - **Cutover**—This option is only displayed if a migration is in a waiting to cutover state. Select this option to begin cutover.
  - **Delete**—Select this option to delete the migration. You will see an additional option to **Delete cloud resources created by selected migrations**. Be careful when deleting the cloud resources. Only delete them if your migration failed or you do not want to use the

target. Deleting the cloud resources will delete everything in the target cloud, including your target server.

- **Migration status**—The **Status** column has different colors and icons to indicate the health of your migrations.
  - **Green arrows**—A green circle with two curved arrows indicates the migration process is running.
  - **Green checkmark**—A green circle with a white checkmark indicates the migration process is complete.
  - **Yellow exclamation point**—A yellow circle with a white exclamation point indicates a pending or warning status.
  - **Red X**—A red circle with a white X indicates an error status.

# Migrating to Amazon Web Services

Make sure you have reviewed the *Target requirements for Amazon Web Services* on page 16 before beginning your migration.



Due to limitations provisioning storage for Linux servers in the cloud, Linux source servers can only be migrated to custom target environments at this time. This means you will need to have an existing target server already provisioned and running in your cloud to use as a custom target.

1. You can start a migration from several locations.
  - On the **Migrations** page, click **Create**.
  - On the **Servers** page, click **Migrate** from a server's **Actions** menu.
  - On the **Servers** page, select one or more servers and click **Migrate** in the toolbar.
2. Identify your migration settings.

### Create Migration

1

2

3

4

➔

MIGRATION

SOURCE

TARGET

CONFIGURATION

Group Name

Alpha and Beta migration

Target Environment

Amazon Web Services

▼

Target Location

US East (Virginia)

▼

Target Key Pair

No Key Pair

▼

?

Target RDP Access

RDP port open for specific sub-domain

▼

Sub-domain for RDP Access

172.26.0.1/16

☒ Create public IP addresses for target servers

☒ Automatically cutover after synchronization completes

☐ Shut down source when cutover starts

☐ Compress data during migration

☒ Remove Carbonite Migrate from target environment after successful migration

Cancel

Next ➔

- **Group Name**—Specify a unique group name for this migration. Even if you are only migrating one server, a group will be created.
- **Target Environment**—Specify your AWS target environment.
- **Target Location**—Select the location where you want to migrate your source servers.
- **Target Key Pair**—Select an Amazon Web Services EC2 key pair from the list. The key pair selected will be used for public-key cryptography to encrypt and decrypt login information for the target AWS EC2 instance.
- **Target RDP Access**—If desired, you can select to open the RDP port for access to the target after cutover.
  - **No RDP**—Select this option if you do not want to open the RDP port for access to the target after cutover.
  - **RDP port open for specific sub-domain**—Select this option if you want to open the RDP port for specific IP addresses. For **Sub-domain for RDP access**, specify the IP addresses by using standard CIDR notation, for example 172.26.0.1/16.
  - **RDP port open for all IPs**—Select this option if you want to open the RDP port for all IP addresses. This option is not recommended.
- **Options**—Specify your migration options.
  - **Create public IP addresses for target servers**—If you are using a proxy server and have this option selected, Amazon Web Services will create a public IP address on the target. If this option is not selected, Amazon Web Services will not create a public IP address on the target, and you must have a VPN connection so your source and target can communicate. This option is not available if you are not using a proxy server.
  - **Automatically cutover after synchronization completes**—When selected, cutover will begin automatically as soon as synchronization is complete. If this option is not selected, the migration process will continue replicating data changes but will not begin the cutover step until you click **Cutover** to continue the migration process.
  - **Shut down source when cutover starts**—When selected, the source will be shut down automatically when cutover begins. In most cases, you should shut down the source because conflicts can occur if you have two servers with the same name on the same network. Shutting down the source server allows the target to come up without conflicts. If this option is not selected, the source will remain online when cutover begins.
  - **Compress data during migration**—When selected, data will be compressed as it is sent across the wire from the source to the target. Compression can reduce the amount of bandwidth consumed as well as the time that it takes to synchronize all of the data to the target. If this option is not selected, the data is not compressed.
  - **Remove Carbonite Migrate from target environment after successful migration**—When selected, Carbonite Migrate will be uninstalled and any security groups needed for the migration, if created, are removed. This option is not available for custom target environments.



If Amazon Web Services cannot delete the security group because of dependent objects, you will need to manually disassociate or remove the



security group from all dependent objects then manually delete it.

3. Click **Next** to continue.
4. Identify the source servers you want to migrate. Only servers in a ready to migrate status will be displayed. Multiple selected servers will be part of a group migration.

Create Migration

< 1 2 3 4 >

MIGRATION SOURCE TARGET CONFIGURATION

Alpha and Beta migration

2 Selected [Select All](#) [Clear All Selections](#) 20 per page

<input type="checkbox"/>	Name ^	CPU	Platform ↕	Volume Size ↕	Source Route
<input checked="" type="checkbox"/>	ALPHA	1 cores, 4.0 GB memory	Microsoft Windows Server 2008 R2 Enterprise	39.9 GB	Auto
<input checked="" type="checkbox"/>	BETA	1 cores, 4.0 GB memory	Microsoft Windows Server 2008 R2 Enterprise	39.9 GB	Auto
<input type="checkbox"/>	DELTA	2 cores, 4.0 GB memory	Microsoft Windows Server 2012 R2 Datacenter	39.7 GB	172.31.186.63
<input type="checkbox"/>	GAMMA	2 cores, 4.0 GB memory	CentOS	82.1 GB	Auto

[Cancel](#) [Previous](#) [Next](#)

- **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.
- **Search**—Text entered in the search 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.
- **Name**—You can expand a server by clicking on the arrow to the left of the name to see what will be migrated on that server. By default, all volumes are selected for migration. Deselect any volumes you do not want to migrate, but be careful when excluding volumes as they may compromise the integrity of your installed applications on the target. There are some volumes that you will be unable to exclude, because they are required for protection. For example, boot volumes cannot be excluded.
- **CPU, Platform, and Volume Size**—These fields are display only and show the cores, memory, operating system, and total volume size on the source. When a server's individual volumes are displayed, you can see the file system, space remaining, and volume size for each individual volume.

- **Source Route**—If your source environment is Amazon Web Services, vSphere, or vCloud, the source route is automatically determined and you cannot change it. If your source environment is custom or Active Directory, you need to select a source IP address to use for source to target communication. You can also select **Add New IP** to specify a different source IP address for migration.
5. Click **Next** to continue.
  6. For each source you are migrating, identify how you want to create the target server.

Create Migration

1 2 3 4  
MIGRATION SOURCE TARGET CONFIGURATION  
Alpha and Beta migration

Source	Target	Encrypted EBS	VM Size
ALPHA	T-ALPHA	<input type="checkbox"/>	m3.medium (1 core, 3.75 GB memory)
Template Windows_Server-2008-R2_SP1-English-64Bit-Base-2017.05.10 <a href="#">Edit</a> <a href="#">Target Volume</a>			
BETA	T-BETA	<input type="checkbox"/>	m3.medium (1 core, 3.75 GB memory)
Template Windows_Server-2008-R2_SP1-English-64Bit-Base-2017.05.10 <a href="#">Edit</a> <a href="#">Target Volume</a>			

Cancel < Previous Next >

- **Source**—This field is display only and shows the name of the source.
- **Target**—This is the virtual machine display name of the target server that will be created in the cloud. Modify the name of the target, if needed. The name cannot exceed 255 characters.
- **Encrypted EBS**—Enable this checkbox if you want the virtual machine target server to be created in an encrypted Elastic Block Store (EBS) in Amazon. You can enable or disable all of your targets for EBS by using the checkbox in the column heading. Keep in mind, if you are using an Amazon Machine Image with encrypted EBS boot volume capability, all volumes will be encrypted. Otherwise, Amazon will only encrypt data volumes.
- **VM Size**—This is the size of the virtual machine target server that will be created in the cloud. The default size is selected to be as close in configuration as possible to your source. Select a different virtual machine size, if needed. The sizes available in the list will change depending on EBS selection.
- **Template**—For each operating system you are migrating, Carbonite Migrate will select an operating system template that will be used to create the target server in the cloud. The selected template is as close in configuration as possible to the operating system being migrated. If you want to use a different template, click **Edit**.
  - a. Identify the type of template you want to search for by completing the following fields.
    - **Name**—Specify a descriptive keyword to search for. If you enter a specific AMI ID (Amazon Machine Image), none of the other search criteria will be used.
    - **Platform**—Specify a platform to search for.

- **Version**—Specify a version to search for.
  - **Publisher**—Select where you want to search in AWS. If you select AWS Marketplace, make sure you have subscribed to the template before selecting it.
  - **Architecture**—Select the architecture to search for.
  - **Category**—Select to search public or private templates.
- b. After you have defined your search criteria, click **Search**.
  - c. If you want to use a different template, select it from the search results.
  - d. If you want to use this same template for all migrations of this operating system, select the option to replace the current template for all migrations.
  - e. Click **Confirm**.



Once you have selected and saved a new template to use, you will see a **Reset Template** button on the search page. Click this button to go back to the template that Carbonite Migrate originally selected.

The list of templates is determined by the AMIs (Amazon Machine Images) that you have available. Carbonite Migrate can only use AMIs that meet the following characteristics.

- **Image type**—machine
- **Root device type**—ebs
- **Virtualization type**—hvm
- **Block device mapping volume type**— standard or gp2

If you do not have an AMI that meets this criteria, your search list will be empty.

- 
- **Target Volume**—Expand this section to see each volume you are protecting. Select the specific **Volume Type** that you want to use. The types available in the drop-down list will depend on the source volume. For example, system volumes do not have as many compatible volumes types as data volumes. See Amazon documentation for details on each volume type.
7. Click **Next** to continue.
  8. Identify the configuration of the target server.

Create Migration

**Configuration**

Network: vpc-2a6abf4f (GN-AWS)

Subnet: subnet-b132dac6 (172.31.16.0/20)

**Credentials**

Provide credentials on the target that have administrator privileges. These will be replaced by the credentials from the source server after cutover.

Username: The local administrator username will be used

Password: [password field] [eye icon] [help icon]

Buttons: Cancel, < Previous, Create

- **Network**—Select the Amazon Virtual Private Cloud (VPC) network that you want the target server to use.
  - **Subnet**—Select the subnet for the selected network.
  - **Credentials**—Specify a temporary local administrator password for creating the target server. Make sure the password meets the security requirements of your target environment. After cutover, the temporary password will be replaced by credentials from the source.
9. When you are ready to begin the migration, click **Create**.
  10. Track your migration on the **Migrations** page. See *Migrations* on page 64 for details.
  11. Once the synchronization has completed, cutover will begin immediately if you selected the automatic cutover option. If you did not select that option, the migration will pause in a **Waiting to cutover** state. During this pause, data changes on the source will continue to synchronize to the target. When you are ready, click on the server name on the **Migrations** page and select **Cutover**. The status will change to **Cutting over**.
  12. When the cutover status is **Done**, the migration is complete. You can delete the migration by clicking the server name on the **Migrations** page and selecting **Delete**. You will see an additional option to **Delete cloud resources created by selected migrations**. Be careful when deleting the cloud resources. Only delete them if your migration failed or you do not want to use the target. Deleting the cloud resources will delete all of the target cloud resources created by this specific migration, including your target server.



You may need to create or update your network settings after the migration is complete. Be sure and check the configuration of firewall, DNAT, SNAT, NAT, PAT, or other networking settings for connectivity.



# Migrating to Google Cloud Platform



Due to limitations provisioning storage for Linux servers in the cloud, Linux source servers can only be migrated to custom target environments at this time. This means you will need to have an existing target server already provisioned and running in your cloud to use as a custom target.

1. You can start a migration from several locations.
  - On the **Migrations** page, click **Create**.
  - On the **Servers** page, click **Migrate** from a server's **Actions** menu.
  - On the **Servers** page, select one or more servers and click **Migrate** in the toolbar.
2. Identify your migration settings.

Create Migration

1 2 3 4

MIGRATION SOURCE TARGET CONFIGURATION

Group Name

Target Environment

Target Location

Target RDP Access

☐ Automatically cutover after synchronization completes

☒ Shut down source when cutover starts

☒ Compress data during migration

☒ Remove Carbonite Migrate from target environment after successful migration

- **Group Name**—Specify a unique group name for this migration. Even if you are only migrating one server, a group will be created.
- **Target Environment**—Specify your Google Cloud Platform target environment.
- **Target Location**—Select the location where you want to migrate your source servers. The list is a combination of available regions and zones.
- **Target RDP Access**—If desired, you can select to open the RDP port for access to the target after cutover.
  - **No RDP**—Select this option if you do not want to open the RDP port for access to the target after cutover.

- **RDP port open for specific sub-domain**—Select this option if you want to open the RDP port for specific IP addresses. For **Sub-domain for RDP access**, specify the IP addresses by using standard CIDR notation, for example 172.26.0.1/16.
- **RDP port open for all IPs**—Select this option if you want to open the RDP port for all IP addresses. This option is not recommended.
- **Options**—Specify your migration options.
  - **Automatically cutover after synchronization completes**—When selected, cutover will begin automatically as soon as synchronization is complete. If this option is not selected, the migration process will continue replicating data changes but will not begin the cutover step until you click **Cutover** to continue the migration process.
  - **Shut down source when cutover starts**—When selected, the source will be shut down automatically when cutover begins. In most cases, you should shut down the source because conflicts can occur if you have two servers with the same name on the same network. Shutting down the source server allows the target to come up without conflicts. If this option is not selected, the source will remain online when cutover begins.
  - **Compress data during migration**—When selected, data will be compressed as it is sent across the wire from the source to the target. Compression can reduce the amount of bandwidth consumed as well as the time that it takes to synchronize all of the data to the target. If this option is not selected, the data is not compressed.
  - **Remove Carbonite Migrate from target environment after successful migration**—When selected, Carbonite Migrate will be uninstalled and any security groups needed for the migration, if created, are removed. This option is not available for custom target environments.

3. Click **Next** to continue.

4. Identify the source servers you want to migrate. Only servers in a ready to migrate status will be displayed. Multiple selected servers will be part of a group migration.

Create Migration

2 Selected [Select All](#) [Clear All Selections](#) 20 per page

<input type="checkbox"/>	Name ^	CPU	Platform ^	Volume Size ^	Source Route
<input checked="" type="checkbox"/>	ALPHA	1 cores, 4.0 GB memory	Microsoft Windows Server 2008 R2 Enterprise	39.9 GB	Auto
<input checked="" type="checkbox"/>	BETA	1 cores, 4.0 GB memory	Microsoft Windows Server 2008 R2 Enterprise	39.9 GB	Auto
<input type="checkbox"/>	DELTA	2 cores, 4.0 GB memory	Microsoft Windows Server 2012 R2 Datacenter	39.7 GB	172.31.186.63
<input type="checkbox"/>	GAMMA	2 cores, 4.0 GB memory	CentOS	82.1 GB	Auto

[Cancel](#) [Previous](#) [Next](#)

- **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.
  - **Search**—Text entered in the search 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.
  - **Name**—You can expand a server by clicking on the arrow to the left of the name to see what will be migrated on that server. By default, all volumes are selected for migration. Deselect any volumes you do not want to migrate, but be careful when excluding volumes as they may compromise the integrity of your installed applications on the target. There are some volumes that you will be unable to exclude, because they are required for protection. For example, boot volumes cannot be excluded.
  - **CPU, Platform, and Volume Size**—These fields are display only and show the cores, memory, operating system, and total volume size on the source. When a server's individual volumes are displayed, you can see the file system, space remaining, and volume size for each individual volume.
  - **Source Route**—If your source environment is Amazon Web Services, vSphere, or vCloud, the source route is automatically determined and you cannot change it. If your source environment is custom or Active Directory, you need to select a source IP address to use for source to target communication. You can also select **Add New IP** to specify a different source IP address for migration.
5. Click **Next** to continue.
  6. For each source you are migrating, identify how you want to create the target server.

Create Migration

< 1 2 3 4 >  
 MIGRATION SOURCE TARGET CONFIGURATION  
 Alpha and Beta migration

Source	Target	VM Size
Alpha	t-alpha	e2-medium (2 cores, 4 GB memory)
Template windows-server-2012-r2-dc-v20210309		
> Target Volume		
Beta	t-beta	e2-medium (2 cores, 4 GB memory)
Template windows-server-2012-r2-dc-v20210309		
> Target Volume		

Cancel < Previous Next >

- **Source**—This field is display only and shows the name of the source.
- **Target**—This is the virtual machine display name of the target server that will be created in the cloud. Modify the name of the target, if needed. The name cannot exceed 63 characters, and it can only contain lowercase letters, numbers, and dashes (-). The first character must be a lowercase letter and the last character cannot be a dash.
- **VM Size**—This is the size of the virtual machine target server that will be created in the cloud. The default size is selected to be as close in configuration as possible to your source. Select a different virtual machine size, if needed.
- **Template**—For each operating system you are migrating, Carbonite Migrate will attempt to select an operating system template that will be used to create the target server in the cloud. The selection will depend on if there is a publicly available template that matches the source operating system. Any selected template is as close in configuration as possible to the operating system being migrated.
  - a. If you need to select a template or want to use a different template, click **Edit**.
  - b. If you do not want to see all public templates, narrow the list of templates by completing any of the following fields.
    - **Name**—Specify a descriptive keyword to search for. This field will search for matches in both template name and template description.
    - **Project**—Specify a project name to search.
    - **Show Deprecated**—Enable this option if you want to include deprecated (legacy) templates.
  - a. After you have defined your search criteria, click **Search**.
  - b. If you want to use a different template, select it from the search results.
  - c. If you want to use this same template for all migrations of this operating system, select the option to replace the current template for all migrations.
  - d. Click **Confirm**.



Once you have selected and saved a new template to use, you will see a **Reset Template** button on the search page. Click this button to go back to the template that Carbonite Migrate originally selected.

---

- **Target Volume**—Expand this section to see each volume you are protecting. Select the specific **Volume Type** that you want to use. The types available in the drop-down list will depend on the source volume and what you have available in your project. See Google documentation for details on each volume type.
7. Click **Next** to continue.
  8. Identify the configuration of the target server.

Create Migration

Configuration



Network

Subnet

Credentials

Provide credentials on the target that have administrator privileges. These will be replaced by the credentials from the source server after cutover.

Username The local administrator username will be used

Password   

- **Network**—Select the virtual private network that you have available to your project that want the target server to use.
  - **Subnet**—Select the subnet for the selected network.
  - **Credentials**—Specify a temporary local administrator password for creating the target server. The password must contain lowercase and uppercase letters and at least one number. The password must be at least eight characters. After cutover, the temporary password will be replaced by credentials from the source.
9. When you are ready to begin the migration, click **Create**.
  10. Track your migration on the **Migrations** page. See *Migrations* on page 64 for details.
  11. Once the synchronization has completed, cutover will begin immediately if you selected the automatic cutover option. If you did not select that option, the migration will pause in a **Waiting to cutover** state. During this pause, data changes on the source will continue to synchronize to the target. When you are ready, click on the server name on the **Migrations** page and select **Cutover**. The status will change to **Cutting over**.
  12. When the cutover status is **Done**, the migration is complete. You can delete the migration by clicking the server name on the **Migrations** page and selecting **Delete**. You will see an additional option to **Delete cloud resources created by selected migrations**. Be careful when deleting the cloud resources. Only delete them if your migration failed or you do not want to use the target. Deleting the cloud resources will delete all of the target cloud resources created by this specific migration, including your target server.



You may need to create or update your network settings after the migration is complete. Be sure and check the configuration of firewall, DNAT, SNAT, NAT, PAT, or other networking settings for connectivity.

---

# Migrating to Microsoft Azure Resource Manager

Make sure you have reviewed the *Target requirements for Microsoft Azure Resource Manager* on page 18 before beginning your migration.



Due to limitations provisioning storage for Linux servers in the cloud, Linux source servers can only be migrated to custom target environments at this time. This means you will need to have an existing target server already provisioned and running in your cloud to use as a custom target.

1. You can start a migration from several locations.
  - On the **Migrations** page, click **Create**.
  - On the **Servers** page, click **Migrate** from a server's **Actions** menu.
  - On the **Servers** page, select one or more servers and click **Migrate** in the toolbar.
2. Identify your migration settings.

### Create Migration

1

2

3

4

MIGRATION

SOURCE

TARGET

CONFIGURATION

Group Name

Alpha and Beta migration

Target Environment

Microsoft Azure Resource Manager

Target Location

East US

Target RDP Access

RDP port open for specific sub-domain

Sub-domain for RDP Access

172.26.0.1/16

☒ Automatically cutover after synchronization completes

☐ Shut down source when cutover starts

☒ Compress data during migration

☒ Remove Carbonite Migrate from target environment after successful migration

Cancel

Next >

- **Group Name**—Specify a unique group name for this migration. Even if you are only migrating one server, a group will be created.
- **Target Environment**—Specify your Azure Resource Manager target environment.
- **Target Location**—Select the location where you want to migrate your source servers.

- **Target RDP Access**—If desired, you can select to open the RDP port for access to the target after cutover.
  - **No RDP**—Select this option if you do not want to open the RDP port for access to the target after cutover.
  - **RDP port open for specific sub-domain**—Select this option if you want to open the RDP port for specific IP addresses. For **Sub-domain for RDP access**, specify the IP addresses by using standard CIDR notation, for example 172.26.0.1/16.
  - **RDP port open for all IPs**—Select this option if you want to open the RDP port for all IP addresses. This option is not recommended.
- **Options**—Specify your migration options.
  - **Automatically cutover after synchronization completes**—When selected, cutover will begin automatically as soon as synchronization is complete. If this option is not selected, the migration process will continue replicating data changes but will not begin the cutover step until you click **Cutover** to continue the migration process.
  - **Shut down source when cutover starts**—When selected, the source will be shut down automatically when cutover begins. In most cases, you should shut down the source because conflicts can occur if you have two servers with the same name on the same network. Shutting down the source server allows the target to come up without conflicts. If this option is not selected, the source will remain online when cutover begins.
  - **Compress data during migration**—When selected, data will be compressed as it is sent across the wire from the source to the target. Compression can reduce the amount of bandwidth consumed as well as the time that it takes to synchronize all of the data to the target. If this option is not selected, the data is not compressed.
  - **Remove Carbonite Migrate from target environment after successful migration**—When selected, Carbonite Migrate will be uninstalled and any security groups needed for the migration, if created, are removed. This option is not available for custom target environments.

3. Click **Next** to continue.

4. Identify the source servers you want to migrate. Only servers in a ready to migrate status will be displayed. Multiple selected servers will be part of a group migration.

Create Migration

<
1
2
3
4
>

MIGRATION  
SOURCE    TARGET    CONFIGURATION

Alpha and Beta migration

2 Selected   [Select All](#)   [Clear All Selections](#)   20 per page  

<input type="checkbox"/>	Name ^	CPU	Platform ↕	Volume Size ↕	Source Route
<input checked="" type="checkbox"/>	ALPHA	1 cores, 4.0 GB memory	Microsoft Windows Server 2008 R2 Enterprise	39.9 GB	Auto
<input checked="" type="checkbox"/>	BETA	1 cores, 4.0 GB memory	Microsoft Windows Server 2008 R2 Enterprise	39.9 GB	Auto
<input type="checkbox"/>	DELTA	2 cores, 4.0 GB memory	Microsoft Windows Server 2012 R2 Datacenter	39.7 GB	172.31.186.63
<input type="checkbox"/>	GAMMA	2 cores, 4.0 GB memory	CentOS	82.1 GB	Auto



- **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.
  - **Search**—Text entered in the search 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.
  - **Name**—You can expand a server by clicking on the arrow to the left of the name to see what will be migrated on that server. By default, all volumes are selected for migration. Deselect any volumes you do not want to migrate, but be careful when excluding volumes as they may compromise the integrity of your installed applications on the target. There are some volumes that you will be unable to exclude, because they are required for protection. For example, boot volumes cannot be excluded.
  - **CPU, Platform, and Volume Size**—These fields are display only and show the cores, memory, operating system, and total volume size on the source. When a server's individual volumes are displayed, you can see the file system, space remaining, and volume size for each individual volume.
  - **Source Route**—If your source environment is Amazon Web Services, vSphere, or vCloud, the source route is automatically determined and you cannot change it. If your source environment is custom or Active Directory, you need to select a source IP address to use for source to target communication. You can also select **Add New IP** to specify a different source IP address for migration.
5. Click **Next** to continue.
  6. For each source you are migrating, identify how you want to create the target server.

Create Migration

Resource Group

Res-grp

Migrate Project

MigProj

Source	Target	VM Size
Alpha	T-Alpha	Standard_DS1_v2 (1 core, 3.5 GB memory)
Template WindowsServer 2012-R2-Datacenter		
Beta	T-Beta	Standard_DS1_v2 (1 core, 3.5 GB memory)
Template WindowsServer 2012-R2-Datacenter		

Cancel Previous Next

- **Resource Group**—Select a resource group which contains the Azure migrate project. The resource groups listed are dependent on the subscription used to create the target environment. You can also select **Automatically generate** to have a resource group created for you.
- **Migrate Project**—Select the Azure migrate project that you want to use for the migration. This is the migrate project under Azure Migrate in the Azure portal. A migrate project is used to store discovery, assessment, and migration metadata reported by the source environment. If you do not have a migrate project for the **Resource Group** in the selected **Target Location**, select a different **Resource Group** or go back and select a different **Target Location** (that contains a **Resource Group** with a migrate project). If you do not have any migrate projects, go back and use the **Microsoft Azure Resource Manager (legacy) Target Environment** type.
- **Source**—This field is display only and shows the name of the source.
- **Target**—This is the virtual machine display name of the target server that will be created in the cloud. Modify the name of the target, if needed. Since the same name will be used for the target hostname during the migration process, the target name cannot be identical to the source name. After the migration is complete, the target hostname will be updated to match the original source name. Due to Azure Resource Manager limitations, the name cannot be longer than 15 characters. The name cannot contain special characters `/'"[]:|<>+=,;?*@&` or begin with `'_'` or end with `'.'`.
- **VM Size**—This value shows the virtual machines as offered by Azure. Select a different virtual machine size, if needed. Virtual machines with less than 1 GB of memory will automatically be excluded from the available list.

- **Template**—For each operating system you are migrating, Carbonite Migrate will select an operating system template that will be used to create the target server in the cloud. The selected template is as close in configuration as possible to the operating system being migrated. If you want to use a different template, click **Edit**.
    - a. Identify the type of template you want to search for by completing the following fields. At a minimum, you need to specify the publisher.
      - **Name**—Specify a descriptive keyword to search for.
      - **Publisher**—Specify the publisher name to search for. This is a required search parameter.
      - **Offer**—Specify the offer to search for. The offer is a descriptor for the operating system family.
      - **SKU**—Specify the SKU to search for. The SKU is a descriptor for the release or version of the offer.
    - b. After you have defined your search criteria, click **Search**.
    - c. If you want to use a different template, select it from the search results.
    - d. If you want to use this same template for all migrations of this operating system, select the option to replace the current template for all migrations.
    - e. Click **Confirm**.
- 



Once you have selected and saved a new template to use, you will see a **Reset Template** button on the search page. Click this button to go back to the template that Carbonite Migrate originally selected.

---

7. Click **Next** to continue.
8. Identify the configuration of the target server.

Create Migration

1 MIGRATION  
Alpha and Beta migration

2 SOURCE

3 TARGET

4 CONFIGURATION

Configuration

Network



Subnet

Storage Account

Credentials

Provide credentials on the target that have administrator privileges. These will be replaced by the credentials from the source server after cutover.

Username

Password   

- **Network**—Select the network that you want the target server to use. If your migration requires a VPN, you must select a network that has a VPN connection associated to it.
  - **Subnet**—Select the subnet for the selected network.
  - **Storage account**—Select a storage account to use when creating the target server in the cloud. You can also select **Automatically generate** to have the storage account created for you.
  - **Credentials**—Specify a temporary user and password for creating the target server. Make sure the password meets the security requirements of your target environment. After cutover, these credentials will be replaced by credentials from the source.
9. When you are ready to begin the migration, click **Create**.
  10. Track your migration on the **Migrations** page. See *Migrations* on page 64 for details.
  11. Once the synchronization has completed, cutover will begin immediately if you selected the automatic cutover option. If you did not select that option, the migration will pause in a **Waiting to cutover** state. During this pause, data changes on the source will continue to synchronize to the target. When you are ready, click on the server name on the **Migrations** page and select **Cutover**. The status will change to **Cutting over**.
  12. When the cutover status is **Done**, the migration is complete. You can delete the migration by clicking the server name on the **Migrations** page and selecting **Delete**. You will see an additional option to **Delete cloud resources created by selected migrations**. Be careful when deleting the cloud resources. Only delete them if your migration failed or you do not want to use the target. Deleting the cloud resources will delete all of the target cloud resources created by this specific migration, including your target server.



You may need to create or update your network settings after the migration is complete. Be sure and check the configuration of firewall, DNAT, SNAT, NAT, PAT, or other networking settings for connectivity.

---

# Migrating to Microsoft Azure Resource Manager (legacy)

Make sure you have reviewed the *Target requirements for Microsoft Azure Resource Manager (legacy)* on page 24 before beginning your migration.



Due to limitations provisioning storage for Linux servers in the cloud, Linux source servers can only be migrated to custom target environments at this time. This means you will need to have an existing target server already provisioned and running in your cloud to use as a custom target.

1. You can start a migration from several locations.
  - On the **Migrations** page, click **Create**.
  - On the **Servers** page, click **Migrate** from a server's **Actions** menu.
  - On the **Servers** page, select one or more servers and click **Migrate** in the toolbar.
2. Identify your migration settings.

### Create Migration

1

2

3

4

➔

MIGRATION

SOURCE

TARGET

CONFIGURATION

Group Name

Alpha and Beta migration

Target Environment

Microsoft Azure Resource Manager (legacy) 

▼

Target Location

East US 

▼

Target RDP Access

RDP port open for specific sub-domain 

▼

Sub-domain for RDP Access

172.26.0.1/16

☒ Automatically cutover after synchronization completes

☐ Shut down source when cutover starts

☒ Compress data during migration

☒ Remove Carbonite Migrate from target environment after successful migration

Cancel

Next ➔

- **Group Name**—Specify a unique group name for this migration. Even if you are only migrating one server, a group will be created.
- **Target Environment**—Specify your Azure Resource Manager (legacy) target environment.
- **Target Location**—Select the location where you want to migrate your source servers.

- **Target RDP Access**—If desired, you can select to open the RDP port for access to the target after cutover.
  - **No RDP**—Select this option if you do not want to open the RDP port for access to the target after cutover.
  - **RDP port open for specific sub-domain**—Select this option if you want to open the RDP port for specific IP addresses. For **Sub-domain for RDP access**, specify the IP addresses by using standard CIDR notation, for example 172.26.0.1/16.
  - **RDP port open for all IPs**—Select this option if you want to open the RDP port for all IP addresses. This option is not recommended.
- **Options**—Specify your migration options.
  - **Automatically cutover after synchronization completes**—When selected, cutover will begin automatically as soon as synchronization is complete. If this option is not selected, the migration process will continue replicating data changes but will not begin the cutover step until you click **Cutover** to continue the migration process.
  - **Shut down source when cutover starts**—When selected, the source will be shut down automatically when cutover begins. In most cases, you should shut down the source because conflicts can occur if you have two servers with the same name on the same network. Shutting down the source server allows the target to come up without conflicts. If this option is not selected, the source will remain online when cutover begins.
  - **Compress data during migration**—When selected, data will be compressed as it is sent across the wire from the source to the target. Compression can reduce the amount of bandwidth consumed as well as the time that it takes to synchronize all of the data to the target. If this option is not selected, the data is not compressed.
  - **Remove Carbonite Migrate from target environment after successful migration**—When selected, Carbonite Migrate will be uninstalled and any security groups needed for the migration, if created, are removed. This option is not available for custom target environments.



If you are using Amazon Web Services Administrator credentials and are migrating Windows 2008 R2 or Windows 2016 to Microsoft Azure Classic or Microsoft Azure Resource Manager, Carbonite Migrate may not be uninstalled during the cleanup process. In this case, you should not select the option to remove Carbonite Migrate after successful migration. To complete the uninstall, you need to remove the registry key HKEY\_CLASSES\_ROOT\Installer\Products\F64ABE42F8D664043BBEEE7E95CA9077\Transform from the target server after cutover, and then you can manually uninstall Carbonite Migrate.

---

3. Click **Next** to continue.
4. Identify the source servers you want to migrate. Only servers in a ready to migrate status will be displayed. Multiple selected servers will be part of a group migration.

Create Migration

2 Selected [Select All](#) [Clear All Selections](#) 20 per page

<input type="checkbox"/>	Name ^	CPU	Platform ↕	Volume Size ↕	Source Route
<input checked="" type="checkbox"/>	ALPHA	1 cores, 4.0 GB memory	Microsoft Windows Server 2008 R2 Enterprise	39.9 GB	Auto
<input checked="" type="checkbox"/>	BETA	1 cores, 4.0 GB memory	Microsoft Windows Server 2008 R2 Enterprise	39.9 GB	Auto
<input type="checkbox"/>	DELTA	2 cores, 4.0 GB memory	Microsoft Windows Server 2012 R2 Datacenter	39.7 GB	172.31.186.63
<input type="checkbox"/>	GAMMA	2 cores, 4.0 GB memory	CentOS	82.1 GB	Auto

[Cancel](#) [Previous](#) [Next](#)

- **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.
  - **Search**—Text entered in the search 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.
  - **Name**—You can expand a server by clicking on the arrow to the left of the name to see what will be migrated on that server. By default, all volumes are selected for migration. Deselect any volumes you do not want to migrate, but be careful when excluding volumes as they may compromise the integrity of your installed applications on the target. There are some volumes that you will be unable to exclude, because they are required for protection. For example, boot volumes cannot be excluded.
  - **CPU, Platform, and Volume Size**—These fields are display only and show the cores, memory, operating system, and total volume size on the source. When a server's individual volumes are displayed, you can see the file system, space remaining, and volume size for each individual volume.
  - **Source Route**—If your source environment is Amazon Web Services, vSphere, or vCloud, the source route is automatically determined and you cannot change it. If your source environment is custom or Active Directory, you need to select a source IP address to use for source to target communication. You can also select **Add New IP** to specify a different source IP address for migration.
5. Click **Next** to continue.
  6. For each source you are migrating, identify how you want to create the target server.



Create Migration

Resource Group

Automatically generate

Source	Target	VM Size
ALPHA	T-ALPHA	Standard_DS1_v2 (1 core, 3.5 GB memory)
Template 2.127.20170510 <a href="#">Edit</a>		
BETA	T-BETA	Standard_DS1_v2 (1 core, 3.5 GB memory)
Template 2.127.20170510 <a href="#">Edit</a>		

Cancel < Previous Next >

- **Resource Group**—Select a resource group where you want the target server to be located in the cloud. The resource groups listed are dependent on the subscription used to create the target environment. You can also select **Automatically generate** to have a resource group created for you.
- **Source**—This field is display only and shows the name of the source.
- **Target**—This is the virtual machine display name of the target server that will be created in the cloud. Modify the name of the target, if needed. Since the same name will be used for the target hostname during the migration process, the target name cannot be identical to the source name. After the migration is complete, the target hostname will be updated to match the original source name. Due to Azure Resource Manager limitations, the name cannot be longer than 15 characters. The name cannot contain special characters /""[]:|<>+=; ,?\*@& or begin with '\_' or end with '.' or '-'.
- **VM Size**—This value shows the virtual machines as offered by Azure. Select a different virtual machine size, if needed. Virtual machines with less than 1 GB of memory will automatically be excluded from the available list.
- **Template**—For each operating system you are migrating, Carbonite Migrate will select an operating system template that will be used to create the target server in the cloud. The selected template is as close in configuration as possible to the operating system being migrated. If you want to use a different template, click **Edit**.
  - Identify the type of template you want to search for by completing the following fields. At a minimum, you need to specify the publisher.
    - **Name**—Specify a descriptive keyword to search for.
    - **Publisher**—Specify the publisher name to search for. This is a required search parameter.
    - **Offer**—Specify the offer to search for. The offer is a descriptor for the operating system family.
    - **SKU**—Specify the SKU to search for. The SKU is a descriptor for the release or version of the offer.
  - After you have defined your search criteria, click **Search**.

- c. If you want to use a different template, select it from the search results.
- d. If you want to use this same template for all migrations of this operating system, select the option to replace the current template for all migrations.
- e. Click **Confirm**.



Once you have selected and saved a new template to use, you will see a **Reset Template** button on the search page. Click this button to go back to the template that Carbonite Migrate originally selected.

Because the format of the template information varies between vendors, the date of the template may be incorrect based on the available template information.

7. Click **Next** to continue.
8. Identify the configuration of the target server.

Create Migration

< 1 2 3 4

MIGRATION SOURCE TARGET CONFIGURATION

Alpha and Beta migration

Configuration

Network VirtualNetwork

Subnet GatewaySubnet

Storage Account defaultnetworking

Credentials

Provide credentials on the target that have administrator privileges. These will be replaced by the credentials from the source server after cutover.

Username administrator

Password .....

Cancel < Previous Create

- **Network**—Select the network that you want the target server to use. If your migration requires a VPN, you must select a network that has a VPN connection associated to it.
  - **Subnet**—Select the subnet for the selected network.
  - **Storage account**—Select a storage account to use when creating the target server in the cloud. You can also select **Automatically generate** to have the storage account created for you.
  - **Credentials**—Specify a temporary user and password for creating the target server. Make sure the password meets the security requirements of your target environment. After cutover, these credentials will be replaced by credentials from the source.
9. When you are ready to begin the migration, click **Create**.

10. Track your migration on the **Migrations** page. See *Migrations* on page 64 for details.
11. Once the synchronization has completed, cutover will begin immediately if you selected the automatic cutover option. If you did not select that option, the migration will pause in a **Waiting to cutover** state. During this pause, data changes on the source will continue to synchronize to the target. When you are ready, click on the server name on the **Migrations** page and select **Cutover**. The status will change to **Cutting over**.
12. When the cutover status is **Done**, the migration is complete. You can delete the migration by clicking the server name on the **Migrations** page and selecting **Delete**. You will see an additional option to **Delete cloud resources created by selected migrations**. Be careful when deleting the cloud resources. Only delete them if your migration failed or you do not want to use the target. Deleting the cloud resources will delete all of the target cloud resources created by this specific migration, including your target server.



You may need to create or update your network settings after the migration is complete. Be sure and check the configuration of firewall, DNAT, SNAT, NAT, PAT, or other networking settings for connectivity.

---

# Migrating to Microsoft Azure Classic

Make sure you have reviewed the *Target requirements for Microsoft Azure Classic* on page 25 before beginning your migration.



Due to limitations provisioning storage for Linux servers in the cloud, Linux source servers can only be migrated to custom target environments at this time. This means you will need to have an existing target server already provisioned and running in your cloud to use as a custom target.

1. You can start a migration from several locations.
  - On the **Migrations** page, click **Create**.
  - On the **Servers** page, click **Migrate** from a server's **Actions** menu.
  - On the **Servers** page, select one or more servers and click **Migrate** in the toolbar.
2. Identify your migration settings.

### Create Migration

1

2

3

4

>

MIGRATION

SOURCE

TARGET

CONFIGURATION

Group Name

Alpha and Beta migration

Target Environment

Microsoft Azure Classic

▼

Target Location

Virtual-Network

▼

Subnet

Virtual-Network-01 (10.76.76.128/25)

▼

Target RDP Access

RDP port open for specific sub-domain

▼

Sub-domain for RDP Access

172.26.0.1/16

☒ Automatically cutover after synchronization completes

☐ Shut down source when cutover starts

☐ Compress data during migration

☒ Remove Carbonite Migrate from target environment after successful migration

Cancel

Next >

- **Group Name**—Specify a unique group name for this migration. Even if you are only migrating one server, a group will be created.
- **Target Environment**—Specify your Azure target environment.

- **Target Location**—Select the location where you want to migrate your source servers. If your source server is using a proxy, in many, but not all, cases you will need to select a Virtual Network for your **Target Location** and then select a subnet.
- **Subnet**—If you are using a virtual network, you can optionally select a subnet.
- **Target RDP Access**—If desired, you can select to open the RDP port for access to the target after cutover.
  - **No RDP**—Select this option if you do not want to open the RDP port for access to the target after cutover.
  - **RDP port open for specific sub-domain**—Select this option if you want to open the RDP port for specific IP addresses. For **Sub-domain for RDP access**, specify the IP addresses by using standard CIDR notation, for example 172.26.0.1/16.
  - **RDP port open for all IPs**—Select this option if you want to open the RDP port for all IP addresses. This option is not recommended.
- **Options**—Specify your migration options.
  - **Automatically cutover after synchronization completes**—When selected, cutover will begin automatically as soon as synchronization is complete. If this option is not selected, the migration process will continue replicating data changes but will not begin the cutover step until you click **Cutover** to continue the migration process.
  - **Shut down source when cutover starts**—When selected, the source will be shut down automatically when cutover begins. In most cases, you should shut down the source because conflicts can occur if you have two servers with the same name on the same network. Shutting down the source server allows the target to come up without conflicts. If this option is not selected, the source will remain online when cutover begins.
  - **Compress data during migration**—When selected, data will be compressed as it is sent across the wire from the source to the target. Compression can reduce the amount of bandwidth consumed as well as the time that it takes to synchronize all of the data to the target. If this option is not selected, the data is not compressed.
  - **Remove Carbonite Migrate from target environment after successful migration**—When selected, Carbonite Migrate will be uninstalled and any security groups needed for the migration, if created, are removed. This option is not available for custom target environments.



If you are using Amazon Web Services Administrator credentials and are migrating Windows 2008 R2 or Windows 2016 to Microsoft Azure Classic or Microsoft Azure Resource Manager, Carbonite Migrate may not be uninstalled during the cleanup process. In this case, you should not select the option to remove Carbonite Migrate after successful migration. To complete the uninstall, you need to remove the registry key `HKEY_CLASSES_ROOT\Installer\Products\F64ABE42F8D664043BBEEE7E95CA9077\Transform` from the target server after cutover, and then you can manually uninstall Carbonite Migrate.

---

3. Click **Next** to continue.

- Identify the source servers you want to migrate. Only servers in a ready to migrate status will be displayed. Multiple selected servers will be part of a group migration.

Create Migration

1 2 3 4

MIGRATION SOURCE TARGET CONFIGURATION

Alpha and Beta migration

2 Selected Select All Clear All Selections 20 per page Search

<input type="checkbox"/>	Name ^	CPU	Platform ↕	Volume Size ↕	Source Route
<input checked="" type="checkbox"/>	ALPHA	1 cores, 4.0 GB memory	Microsoft Windows Server 2008 R2 Enterprise	39.9 GB	Auto
<input checked="" type="checkbox"/>	BETA	1 cores, 4.0 GB memory	Microsoft Windows Server 2008 R2 Enterprise	39.9 GB	Auto
<input type="checkbox"/>	DELTA	2 cores, 4.0 GB memory	Microsoft Windows Server 2012 R2 Datacenter	39.7 GB	172.31.186.63
<input type="checkbox"/>	GAMMA	2 cores, 4.0 GB memory	CentOS	82.1 GB	Auto

Cancel < Previous Next >

- **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.
  - **Search**—Text entered in the search 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.
  - **Name**—You can expand a server by clicking on the arrow to the left of the name to see what will be migrated on that server. By default, all volumes are selected for migration. Deselect any volumes you do not want to migrate, but be careful when excluding volumes as they may compromise the integrity of your installed applications on the target. There are some volumes that you will be unable to exclude, because they are required for protection. For example, boot volumes cannot be excluded.
  - **CPU, Platform, and Volume Size**—These fields are display only and show the cores, memory, operating system, and total volume size on the source. When a server's individual volumes are displayed, you can see the file system, space remaining, and volume size for each individual volume.
  - **Source Route**—If your source environment is Amazon Web Services, vSphere, or vCloud, the source route is automatically determined and you cannot change it. If your source environment is custom or Active Directory, you need to select a source IP address to use for source to target communication. You can also select **Add New IP** to specify a different source IP address for migration.
- Click **Next** to continue.

6. For each source you are migrating, identify how you want to create the target server.

Create Migration

Source	Target	VM Size	Service Name
ALPHA	T-ALPHA	Standard_D1 (1 core, 3.5 GB memory)	alphasvc
Template a699494373c04fc0bc8f2bb1389d6106__Win2k8R2SP1-Datacenter-20170510-en-us-127GB.vhd			
BETA	T-BETA	Standard_D1 (1 core, 3.5 GB memory)	betasvc
Template a699494373c04fc0bc8f2bb1389d6106__Win2k8R2SP1-Datacenter-20170510-en-us-127GB.vhd			

Cancel Previous Next

- **Source**—This field is display only and shows the name of the source.
- **Target**—This is the virtual machine display name of the target server that will be created in the cloud. Modify the name of the target, if needed. The name cannot contain special characters /"[]:|<>+=,;,\*@& or begin with '\_' or end with '.' or '-'.
- **VM Size**—This value shows the virtual machines as offered by Azure. Select a different virtual machine size, if needed. Virtual machines with less than 1 GB of memory will automatically be excluded from the available list.
- **Service Name**—This value is the name of the service that will be created in Azure before creating the target server. This value must be unique across all of Azure. Modify the service name, if desired.
- **Template**—For each operating system you are migrating, Carbonite Migrate will select an operating system template that will be used to create the target server in the cloud. The selected template is as close in configuration as possible to the operating system being migrated. If you want to use a different template, click **Edit**.
  - Identify the type of template you want to search for by completing the following fields. At a minimum, you need to specify the version and then specify the description or publisher.
    - **Name**—Specify a descriptive keyword to search for. This or **Version** is a required search parameter.
    - **Version**—Specify a version to search for. This is a required search parameter.
    - **Publisher**—Specify the publisher name to search for.
  - After you have defined your search criteria, click **Search**.
  - If you want to use a different template, select it from the search results.
  - If you want to use this same template for all migrations of this operating system, select the option to replace the current template for all migrations.
  - Click **Confirm**.



Once you have selected and saved a new template to use, you will see a **Reset Template** button on the search page. Click this button to go back to the template that Carbonite Migrate originally selected.

Each time you use a Windows template, the size of the system disk created on the target may be different than the size of the system disk on the template based on a calculation. The calculation is the system disk space in use on the source plus minimum amounts of additional free space required by Carbonite Migrate based on your operating system (10 GB for 2008 R2, 14 GB for 2012, and 15 GB for 2012 R2).

- If the calculation is smaller than the template system disk size, the template disk size will be used on the target.
- If the calculation is larger than the template system disk size, the calculation size will be used on the target.
- If the calculation is equal to or larger than the allowed maximum Azure size, the migration will not be allowed.

7. Click **Next** to continue.

8. Identify the configuration of the target server.

Create Migration

< 1 2 3 4

MIGRATION SOURCE TARGET CONFIGURATION

Alpha and Beta migration

Configuration

Storage Account Automatically generate

Credentials

Provide credentials on the target that have administrator privileges. These will be replaced by the credentials from the source server after cutover.

Username administrator

Password .....

Cancel < Previous Create

- **Storage account**—Specify an Azure storage account name. If you do not have a storage account in that location, it will be created automatically.
- **Credentials**—Specify a temporary user and password for creating the target server. Make sure the password meets the security requirements of your target environment. After cutover, these credentials will be replaced by credentials from the source.

9. When you are ready to begin the migration, click **Create**.

10. Track your migration on the **Migrations** page. See *Migrations* on page 64 for details.



11. Once the synchronization has completed, cutover will begin immediately if you selected the automatic cutover option. If you did not select that option, the migration will pause in a **Waiting to cutover** state. During this pause, data changes on the source will continue to synchronize to the target. When you are ready, click on the server name on the **Migrations** page and select **Cutover**. The status will change to **Cutting over**.
  12. When the cutover status is **Done**, the migration is complete. You can delete the migration by clicking the server name on the **Migrations** page and selecting **Delete**. You will see an additional option to **Delete cloud resources created by selected migrations**. Be careful when deleting the cloud resources. Only delete them if your migration failed or you do not want to use the target. Deleting the cloud resources will delete all of the target cloud resources created by this specific migration, including your target server.
- 



You may need to create or update your network settings after the migration is complete. Be sure and check the configuration of firewall, DNAT, SNAT, NAT, PAT, or other networking settings for connectivity.

---

# Migrating to a custom environment

A custom environment as a target environment is only supported when the target server already exists. See *Target requirements for custom environments* on page 26 for complete details on the requirements for using an existing server in a custom target environment.



Due to limitations provisioning storage for Linux servers in the cloud, Linux source servers can only be migrated to custom target environments at this time. This means you will need to have an existing target server already provisioned and running in your cloud to use as a custom target.

1. You can start a migration from several locations.
  - On the **Migrations** page, click **Create**.
  - On the **Servers** page, click **Migrate** from a server's **Actions** menu.
  - On the **Servers** page, select one or more servers and click **Migrate** in the toolbar.
2. Identify your migration settings.

Create Migration

1 2 3 4

MIGRATION SOURCE TARGET CONFIGURATION

Group Name Alpha and Beta migration

Target Environment Custom

Target Location Default location

☐ Automatically cutover after synchronization completes

☒ Shut down source when cutover starts

☒ Compress data during migration

Cancel Next >

- **Group Name**—Specify a unique group name for this migration. Even if you are only migrating one server, a group will be created.
- **Target Environment**—Specify your custom target environment.
- **Target Location**—The location for the existing target you will be selecting will automatically be used.
- **Options**—Specify your migration options.
  - **Automatically cutover after synchronization completes**—When selected, cutover will begin automatically as soon as synchronization is complete. If this option is not selected, the migration process will continue replicating data changes but will not begin the cutover step until you click **Cutover** to continue the migration process.

- **Shut down source when cutover starts**—When selected, the source will be shut down automatically when cutover begins. In most cases, you should shut down the source because conflicts can occur if you have two servers with the same name on the same network. Shutting down the source server allows the target to come up without conflicts. If this option is not selected, the source will remain online when cutover begins.
- **Compress data during migration**—When selected, data will be compressed as it is sent across the wire from the source to the target. Compression can reduce the amount of bandwidth consumed as well as the time that it takes to synchronize all of the data to the target. If this option is not selected, the data is not compressed.

3. Click **Next** to continue.

4. Identify the source servers you want to migrate. Only servers in a ready to migrate status will be displayed. Multiple selected servers will be part of a group migration.

Create Migration

2 Selected [Select All](#) [Clear All Selections](#) 20 per page

<input type="checkbox"/>	Name ^	CPU	Platform ^	Volume Size ^	Source Route
<input checked="" type="checkbox"/>	ALPHA	1 cores, 4.0 GB memory	Microsoft Windows Server 2008 R2 Enterprise	39.9 GB	Auto
<input checked="" type="checkbox"/>	BETA	1 cores, 4.0 GB memory	Microsoft Windows Server 2008 R2 Enterprise	39.9 GB	Auto
<input type="checkbox"/>	DELTA	2 cores, 4.0 GB memory	Microsoft Windows Server 2012 R2 Datacenter	39.7 GB	172.31.186.63
<input type="checkbox"/>	GAMMA	2 cores, 4.0 GB memory	CentOS	82.1 GB	Auto

- **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.
- **Search**—Text entered in the search 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.
- **Name**—You can expand a server by clicking on the arrow to the left of the name to see what will be migrated on that server. By default, all volumes are selected for migration. Deselect any volumes you do not want to migrate, but be careful when excluding volumes as they may compromise the integrity of your installed applications on the target. There are

some volumes that you will be unable to exclude, because they are required for protection. For example, boot volumes cannot be excluded.

- **CPU, Platform, and Volume Size**—These fields are display only and show the cores, memory, operating system, and total volume size on the source. When a server's individual volumes are displayed, you can see the file system, space remaining, and volume size for each individual volume.
- **Source Route**—If your source environment is Amazon Web Services, vSphere, or vCloud, the source route is automatically determined and you cannot change it. If your source environment is custom or Active Directory, you need to select a source IP address to use for source to target communication. You can also select **Add New IP** to specify a different source IP address for migration.

5. Click **Next** to continue.

6. For each source you are migrating, identify the existing target server you want to reuse.

Create Migration

1 2 3 4

MIGRATION SOURCE TARGET CONFIGURATION

Alpha and Beta migration

Source	Target IP Address
ALPHA	172.31.24.200
BETA	172.31.24.201

Cancel < Previous Next >

- **Source**—This field is display only and shows the name of the source.
- **Target**—Specify the IP address of the existing target server you want to use for your migration.

7. Click **Next** to continue.

8. Specify credentials on the target that have administrator privileges. Make sure the password meets the security requirements of your target environment. After cutover, these credentials will be replaced by credentials from the source.

Create Migration

<

1

2

3

4

MIGRATION

SOURCE

TARGET

CONFIGURATION

Alpha and Beta migration

Credentials

Provide credentials on the target that have administrator privileges. These will be replaced by the credentials from the source server after cutover.

Username

administrator

Password

.....

?

Cancel

< Previous

Create

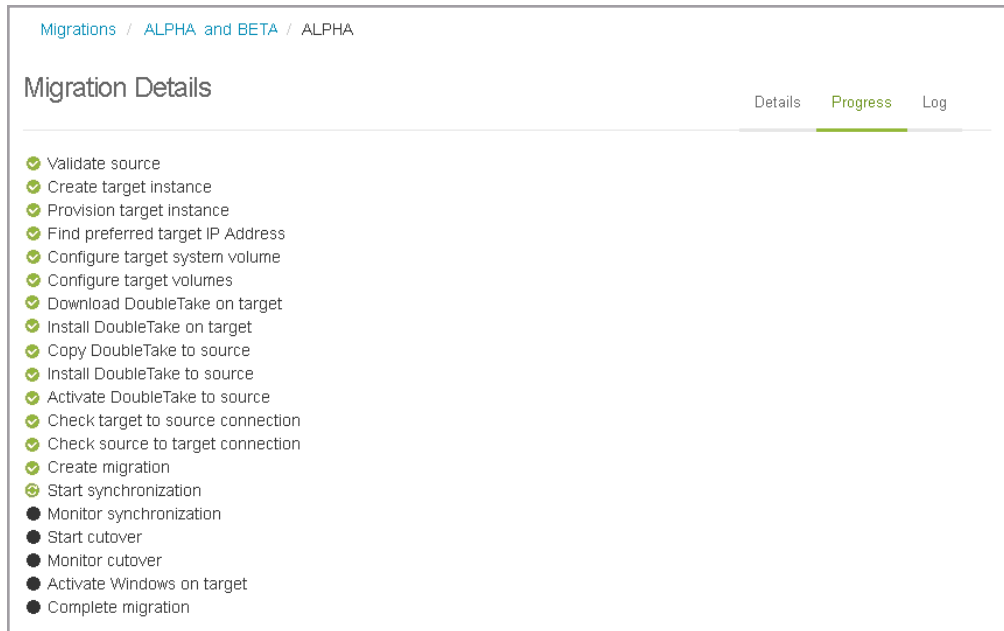
9. When you are ready to begin the migration, click **Create**.
10. Track your migration on the **Migrations** page. See *Migrations* on page 64 for details.
11. Once the synchronization has completed, cutover will begin immediately if you selected the automatic cutover option. If you did not select that option, the migration will pause in a **Waiting to cutover** state. During this pause, data changes on the source will continue to synchronize to the target. When you are ready, click on the server name on the **Migrations** page and select **Cutover**. The status will change to **Cutting over**.
12. When the cutover status is **Done**, the migration is complete. You can delete the migration by clicking the server name on the **Migrations** page and selecting **Delete**. You will see an additional option to **Delete cloud resources created by selected migrations**. Be careful when deleting the cloud resources. Only delete them if your migration failed or you do not want to use the target. Deleting the cloud resources will delete all of the target cloud resources created by this specific migration, including your target server.



You may need to create or update your network settings after the migration is complete. Be sure and check the configuration of firewall, DNAT, SNAT, NAT, PAT, or other networking settings for connectivity.

# Viewing migration details

On the **Migrations** page, click on the name of a migration.



- **Details**—This section displays general overview information about the migration. Click a server name link to see details specific to that server.
- **Progress**—This section shows the workflow of the migration.
  - **Green arrows**—A green circle with two curved arrows indicates that step of the migration process is running.
  - **Green checkmark**—A green circle with a white checkmark indicates that step of the migration completed successfully.
  - **Yellow exclamation point**—A yellow circle with a white exclamation point indicates that step of the migration was in a pending or warning state.
  - **Red X**—A red circle with a white X indicates that step of the migration was in an error state.
  - **Black circle**—A black circle indicates the migration never reached that step.
- **Log**—This section displays a log of the workflow steps.

## Viewing migration history

On the **Migrations** page, click the **History** button. On the **Migration History** page, you will find high-level information showing the history of your migrations. You can sort the migration history list by clicking on any column heading.

Migrations History			
		20 per page	Search
Source ↕	Target ↕	Last Status ↕	Archived ↕
ALPHA	ALPHA	Done	5/10/17 11:31 AM
BETA	BETA	Done	5/10/17 9:08 PM

- **Toolbar and table controls**—The following controls are available on the toolbar and for the table on the **Migration History** page.
  - **per page**—Select the number of table rows to display per page.
  - **Search**—Text entered in the search box will narrow the list displayed to only those rows that contain the search text.
- **Hyperlink control**—Click the name of a migration to see the details for that migration.
  - **Details**—This section contains general information about the migration.
  - **Last Progress**—This section shows the workflow of the migration.
    - **Green arrows**—A green circle with two curved arrows indicates that step of the migration process is running.
    - **Green checkmark**—A green circle with a white checkmark indicates that step of the migration completed successfully.
    - **Yellow exclamation point**—A yellow circle with a white exclamation point indicates that step of the migration was in a pending or warning state.
    - **Red X**—A red circle with a white X indicates that step of the migration was in an error state.
    - **Black circle**—A black circle indicates the migration never reached that step.

---

## Chapter 8 Email notification

You can receive email messages for the alerts generated by Carbonite Migrate.

1. Click on your sign in name, in the upper right corner of the Carbonite Migrate web page and select **My Account**.
2. Select the **Preferences** section.
3. In the **Email** section, specify your email prefix and alert settings.

The screenshot shows the 'My Account' page with the 'Preferences' tab selected. Under the 'Email' section, there is a 'Prefix' field containing 'Carbonite Migrate Alert' and a table for 'Alerts'.

	OK	Warning	Error
Licenses	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Migrations	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Proxies	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Users	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>


- **Prefix**—By default, the subject line of email alerts sent to your account email will be prefaced with Carbonite Migrate Alert. This prefix allows you to recognize and filter emails specific to Carbonite Migrate. You can change or remove the prefix as desired.
  - **Alerts**—Select the type and level of alerts that you want to receive as email messages. If you do not select any type or level, you will not receive alerts as email messages. You will still get alerts in the Carbonite Migrate whether email alerts are enabled or disabled.
4. Click **Save**.



## Chapter 9 Account management

You can manage your Carbonite Migrate account information by clicking on your sign in name, in the upper right corner of the Carbonite Migrate web page and selecting **My Account**.

- **Edit profile information**—You can modify your profile and contact information.
- **Delete account**—The trash icon next to the **My Account** heading allows you to delete your Carbonite Migrate account. You will receive an email requiring you to confirm the deletion.
- **Change account credentials**—You can modify the email address and password used to access the Carbonite Migrate service. This is also the email address that alerts will be sent to, if you have enabled alerts.
- **Two-Step Verification**—You can enable multi-factor authentication to provide more secure access to your 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.

My Account 

Details Preferences Activity Log

First Name

User

Last Name

Name

Company Name

Company Name

Phone Number

317-867-5309

Country

United States

State/Province

Florida


Save

Cancel

Credentials


Email

UserName@domain.com



Password

\*\*\*\*\*



Security

Two-Step Verification

Use your phone to secure your account.

[Get Started](#)