

*Carbonite Availability and  
Carbonite Migrate*

*Replication Reporter User's  
Guide*



## Notices

Carbonite Availability and Carbonite Migrate Replication Reporter User's Guide, version 1.0.0, Tuesday, April 12, 2022

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

Man pages are installed and available on Carbonite Availability and Carbonite Migrate Linux servers. These documents are bound by the same 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; 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 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.

© 2022 Carbonite. All rights reserved.

# Contents

---

- Chapter 1 Overview ..... 4
- Chapter 2 Requirements ..... 5
- Chapter 3 Installing Replication Reporter ..... 6
  - Upgrade notes ..... 6
- Chapter 4 Managing servers for data collection ..... 7
- Chapter 5 Managing reports ..... 8
  - Database fields ..... 11

---

## Chapter 1 Overview

The Carbonite Availability and Carbonite Migrate Replication Reporter is a centralized reporting and analysis tool that allows you to view basic reports and create your own custom detailed reports for any Carbonite Availability and Carbonite Migrate servers in your environment. It can be used to show the overall effectiveness of protection over time and to analyze trends in a data protection scheme. After you install and configure the Replication Reporter, it monitors and collects data from the Carbonite Availability and Carbonite Migrate servers you specify. The collected data is used to generate pre-defined reports, or you can create customized reports.

---

## Chapter 2 Requirements

To use Carbonite Availability and Carbonite Migrate Replication Reporter, your environment must meet the following requirements.

- The Replication Reporter must be installed on Windows 2012 or later.
- The Windows operating system where Replication Reporter is installed must have .NET Core Runtime 3.1.9 or later. If you do not have this, it will be installed when you install Replication Reporter.
- Replication Reporter will install and configure its own database.
- The Replication Reporter can be installed on a Carbonite Availability and Carbonite Migrate source or target server or on an independent server.
- If you are using a firewall, make sure it does not block traffic to and from the Carbonite Availability and Carbonite Migrate source and target servers you want to report on.



Replication Reporter does not support Internet Explorer.

---

---

## Chapter 3 Installing Replication Reporter

1. Close any open applications.
2. Launch the Carbonite Availability and Carbonite Migrate Replication Reporter .exe installation file.



If you are installing on Server Core, copy the .exe file to the Server Core machine using a UNC share, and then launch the installation program from the Server Core machine. The installation screens will display on the Server Core machine.

---

3. At the Welcome page, click **Next** to continue.
4. Review and accept the license agreement to continue with the installation program. Click **Next** to continue.
5. Select the protocol and port to use to access the Replication Reporter. If you select both protocols, the desktop icon will default to the secure protocol.
6. Click **Next** to continue.
7. If you are satisfied with the selections you have made and are ready to begin copying the files, click **Install**.
8. After the files have completed copying, click **Finish** to exit the installation program.

After the installation is complete, you can launch the Replication Reporter interface by using the shortcut icon on the desktop or by entering <https://localhost> or <http://localhost> in a web browser.

## Upgrade notes

Because the Replication Reporter has been completely rewritten, you cannot upgrade if you are using version 8.4.1.46 or earlier. You will need to start fresh with the new version of Replication Reporter. Keep in mind the following if you are upgrading from 8.4.1.46 or earlier.

- You can keep your existing database and data and continue to use it with your existing Replication Reporter version, but you cannot migrate the database or data to the new version.
- The old and new versions can coexist and run on the same server, but you may have to make modifications in how you start the two versions of Replication Reporter on the server.
- If you are using your existing version, you will continue to manage the servers for data collection using the Carbonite Replication Console.
- If you are using the new version, managing the servers for data collection is now completed using the Replication Reporter console.

---

## Chapter 4 Managing servers for data collection

1. Launch the Replication Reporter interface by using the shortcut icon on the desktop or by entering `https://localhost` or `http://localhost` in a web browser.
2. Click the **Servers** link in the header at the top of the page to manage the servers for data collection. If you have no servers you are collecting from, you will also see a **Click here to add a server** link on the **Home** page which takes you to the **Servers** page.
3. Manage the server as needed.
  - **Add a server**—Use the following steps to add a new server for data collection.
    - a. Click the **Add Server** link.
    - b. Specify the server for data collection.
      - **Name**—This is the name or IP address of the Carbonite Availability and Carbonite Migrate server for data collection.
      - **Domain**—If you are working in a domain environment, specify the **Domain**.
      - **User name**—Specify a user that is a member of the **Double-Take Admin** security group on the specified server.
      - **Password**—Specify the password associated with the **User name** you entered.
      - **Collection interval**—Specify the amount of time to wait between data collections.
      - **Retention interval in days**—Specify how long to retain the collected data in the database. Data older than the specified number of days will be deleted.
    - c. Click **Save**.Repeat these steps for all of the servers you want to collect data for.
  - **Edit a server**—Click **Edit** for the server you want to modify. Make the necessary changes and click **Save**.
  - **Remove a server**—Click **Edit** for the server you want to remove. Click **Delete**.

---

## Chapter 5 Managing reports

You can view, download, and email default and custom reports on-demand or at scheduled times.

### Default reports

There are three default reports always available to provide information on jobs, failover activity, and servers. You can find these default reports on the **Home** page.

- **Job Health**—This report shows the most recent status of each job in your database. Keep in mind the following about the job health report.
  - Data is for the time reported as the **Last Successful Collection**. The Replication Reporter is not polling continuously for new data. It polls based on how you configured the server when you added it for monitoring.
  - Job health information is reported by the target of a job, so adding only a source server will not populate this report.
  - If you deleted and re-created a job using the same servers and job name, it may look like you have multiple entries for that job, but what you are seeing is actually a list of unique job IDs.
  - You may see jobs that you have deleted because those jobs will remain in the database (but not updated) until the retention interval you specified for the target server has elapsed.
  - Click a table row to see job health history for the job in that table row. You can see the first and last times that the job was in the specified state.
- **Failover Activity**—This report shows any jobs that have been failed over. These are all failovers until the retention interval you specified for the target server has elapsed. Click a table row to see failover activity history for the failover event in that table row. You can see when failover started and completed.
- **Server Status**—This report shows the most recent status of the servers you are collecting data from. You may see server that you are no longer collecting data from because those servers will remain in the database (but not updated) until the retention interval you specified for the server has elapsed. Click a table row to see server status history for the server in that table row. You can see the first and last times that the server was in the specified state.

### Custom reports

You can create up to five customized reports.

1. Click **Add Custom Report** from any page in the Replication Reporter.
2. Specify a unique **Report Name** that is no more than 20 characters long.
3. Select the fields to include in your report by selecting a field from the **Available Database Fields** list and clicking **Add**.
4. In the **Custom Columns** list, you can modify the column name for the database field, if desired. The name can be no more than 20 characters long.
5. If you want to remove a field from the **Custom Columns** list, highlight the field and click **Remove**.
6. Click **Move Up** and **Move Down** to sort the columns in the desired order.

Once created, the new report will appear as a selectable tab on the **Home** page. You can edit the report or perform any of the actions that are available for the default reports.

### Tile filters

On the **Job Health**, **Server Status**, and any custom reports, you can filter the displayed rows by clicking a tile under the report tab name. Only those table rows that match the tile filter will be displayed. To see all table rows again, click the **All** tile.

### Search

On any of the report tabs, you can filter the displayed rows by entering search criteria. Only those table rows that contain the search criteria will be displayed. To see all table rows again, remove the search criteria.

### Export PDF

You can create a PDF file of the information in a default or custom report.

1. Select one or more rows in a table and click **Export PDF**.
2. Enter a **Report Name** that is no more than 20 characters long.
3. Click **Download**.

The same information that was displayed in the table (for the selected rows) will be included in the PDF file.

### Prepare CSV report

You can create a CSV file of the information in a default or custom report.

1. Select one or more rows in a table and click **Prepare Report**.
2. Enter a **Report Name** that is no more than 20 characters long.
3. Specify how you want to receive the report
  - **Download**—You will be prompted to save or open the file.
  - **Email**—You will be prompted for email information.
    - **Enter Email Information**—Specify the email addresses to send the report to, separating multiple addresses by a comma. Add a subject line for the email that describes the report you are sending. You can also modify the from email address, if desired.
    - **Choose When to Receive Your Report**—Specify when you want the report sent.
      - **Send Now**—The report will be emailed as soon as you click the **Prepare Report** button.
      - **Schedule**—Specify if you want the report to be mailed on a schedule and specify the **Frequency** for the schedule. Make sure you have configured your SMTP server.
4. Click **Prepare Report**.

### Report settings

Click this link, in the header at the top of the page, to view your scheduled reports. You can filter the displayed list by entering search criteria. Only the rows that contain the search criteria will be displayed. If you want to delete a scheduled report, click on that row in the table and click **Delete**. If you want to edit a scheduled report, delete and create a new one.

### **SMTP server settings**

Within the **Report Settings**, click **SMTP Server Settings**. Specify your SMTP configuration and click **Save**.

## Database fields

The table below details the database fields available for custom reports.

---

### Job.JobId

The unique job ID

### Job.Name

The job name

### Job.Timestamp

The date and time the job data was collected

### Job.Type

The job type

### Job.Health

The high level health, or state, of the job. This value equates to the colored icons seen in the first column in the top right pane on the **Jobs** page in the Carbonite Replication Console.

### Job.SourceName

The name of the source server associated with the job

### Job.TargetName

The name of the target server associated with the job

### Job.HighLevelState

The overall state of the job. This value equates to the **Activity** column in the top right pane on the **Jobs** page in the Carbonite Replication Console.

### Job.LowLevelState

The low level state of the job. This value equates to the **Additional Information** field in the bottom right pane on the **Jobs** page in the Carbonite Replication Console.

### Job.TargetState

The state of the data on the target. This value equates to the **Target data state** field in the bottom right pane on the **Jobs** page in the Carbonite Replication Console.

### Connection.TotalBytesSent

The total number of mirror and replication bytes that have been transmitted to the target

### **Connection.RecoveryPointLatency**

The length of time replication is behind on the target compared to the source. This is the time period of replication data that would be lost if a failure were to occur at the current time. This value represents replication data only and does not include mirroring data. If you are mirroring and failover, the data on the target will be at least as far behind as the recovery point latency. It could potentially be further behind depending on the circumstances of the mirror. If mirroring is idle and you failover, the data will only be as far behind as the recovery point latency time.

### **Connection.BandwidthCollar**

The bandwidth limiting that has been set or the keyword **Unlimited** if no bandwidth limit has been set

### **Connection.CompressionEnabled**

Indicates if data is being compressed before it is sent to the target

### **Connection.CompressionLevel**

The level of compression

### **Connection.DiskQueueBytes**

The amount of disk space being used to queue data on the source

### **Connection.InitialMirrorComplete**

Indicates if the initial mirror has been completed

### **Connection.MirrorBytesRemaining**

The total number of mirror bytes that are remaining to be sent from the source to the target.

### **Connection.MirrorBytesSent**

The total number of mirror bytes that have been transmitted to the target

### **Connection.MirrorBytesSkipped**

The total number of bytes that have been skipped when performing a difference mirror. These bytes are skipped because the data is not different on the source and target.

### **Connection.MirrorBytesTransmitted**

The total number of compressed mirror bytes that have been transmitted to the target. If compression is disabled, this will be the same as MirrorBytesSent.

### **Connection.MirrorOpsQueued**

The total number of mirror operations in the queue

### **Connection.MirrorPermillage**

The percentage of the mirror that has been completed

**Connection.MirrorState**

The state of mirroring. This value equates to the **Mirror Status** column in the top right pane on the **Jobs** page in the Carbonite Replication Console.

**Connection.PeerMemoryLow**

Indicates if the target is running low on memory based on the **Amount of system memory to use** setting on the target server's queue properties.

**Connection.ReplicationBytesQueued**

The total number of replication bytes in the source queue

**Connection.ReplicationBytesSent**

The total number of replication bytes that have been transmitted to the target

**Connection.ReplicationBytesTransmitted**

The total number of compressed replication bytes that have been transmitted to the target. If compression is disabled, this will be the same as ReplicationBytesSent.

**Connection.ReplicationOpsQueued**

The total number of replication operations in the queue

**Connection.ReplicationState**

The state of replication. This value equates to the **Replication Status** column in the top right pane on the **Jobs** page in the Carbonite Replication Console.

**Connection.Restoring**

Identifies if the connection is actively restoring

**Connection.SourceAvailable**

Identifies if the target was able to communicate with the source server

**Connection.SourceEngineAvailable**

Identifies if the target was able to communicate with Carbonite Availability and Carbonite Migrate on the source

**Connection.StartTime**

The date and time the connection was initiated

**Connection.TargetAvailable**

Identifies if the source was able to communicate with the target server

**Connection.TargetEngineAvailable**

Identifies if the source was able to communicate with Carbonite Availability and Carbonite Migrate on the target

**Connection.TargetRoute**

The IP address identifying the route to the target

**Connection.TargetQueueBytes**

The number of bytes queued on the target

**Connection.TargetState**

The state of the target

**Connection.TotalBytesTransmitted**

The total number of compressed mirror and replication bytes that have been transmitted to the target. If compression is disabled, this will be the same as TotalBytesSent.

**Connection.TotalOpsQueued**

The total number of mirror and replication operations that are in the source queue

**Connection.TransmissionMode**

Indicates if data is actively being transmitted to the target

**Connection.SourceClusterResourceState**

The state of the Double-Take Source Connection resource, if it is being used by a cluster-aware job on a Carbonite Availability and Carbonite Migrate source cluster

**Connection.CurrentTime**

The date and time the connection data was written to the database

**Connection.MirrorStartTime**

The data and time when the mirroring process started

**Connection.MirrorEndTime**

The data and time when the mirroring process ended

**Connection.RecoveryPointTime**

The date and time replication is synchronized between the source and target. The difference between the TimeStamp and this time is the time period of replication data that would be lost if a failure were to occur at the TimeStamp. This value represents replication data only and does not include mirroring data. If you are mirroring and failover, the data on the target will be at least as far behind as the recovery point time. It could potentially be further behind depending on the circumstances of the mirror. If mirroring is idle and you failover, the data will only be as far behind as the recovery point time.

**Connection.TotalBytesSentCompressed**

The total number of compressed mirror and replication bytes that have been transmitted to the target. If compression is disabled, this statistic will be the same as

**TotalBytesSent.**

**Connection.TotalBytesSentUncompressed**

The total number of mirror and replication bytes that have been transmitted to the target

---