CARBONITE SAFE BACKUP PRO

The Carbonite Safe Server Backup User Guide for 5.x

Table of Contents

System Requirements

The current version of Carbonite Safe Server Backup is designed to support all editions of each Windows operating system listed below*:

- Windows XP (Home or Pro) with SP2 or higher
- Windows Vista

- Windows 7
- Windows Small Business Server 2003
- Windows Server 2003 SP2
- Windows Server 2003 R2
- Windows Small Business Server 2008
- Windows Server 2008
- Windows Server 2008 R2
- Windows Small Business Server 2011
- Windows 8 and 8.1
- Windows 10
- Windows Server 2012
- Windows Server 2012 R2
- Windows Server 2016

**Server core installations are not supported for any version of Windows.*

Hardware Requirements

Your system must meet the following hardware requirements:

- Dual-core CPU
- 2GB of RAM or higher
- 10GB or more of disk space on the drive where CSSB is installed for CSSB's own program and configuration files.
    - Additionally, every disk that contains data to be backed up must have space for VSS snapshots. This space is known as shadowstorage. At least 10% of each disk should be dedicated to shadowstorage. CSSB will automatically create and assign shadowstorage to meet this requirement. You can read more about shadowstorage in this Knowledge Base article.

Other Requirements

Your system must also meet these other requirements:

- Java 8 (or higher) is required for all systems running Windows Vista, Windows Server 2008, and higher.
- Java 7 (or higher) is required for all systems running Windows XP and Windows Server 2003.
- The language selected in the CSSB Installer and the CSSB program must match the Windows system locale.
- The Remote Registry Service must be enabled before installation.
- The Volume Shadow Copy Service (VSS) must be enabled. Although the Volume Shadow Service is enabled by default, it may have been turned off at some point.
    - To ensure that VSS is running:
        - Right-click the **My Computer** (or **Computer**) icon and choose **Manage** from the pop up menu.
        - Expand the *Services and Applications* tree and locate the **Volume Shadow Copy** Service. If it is not started, change the General Properties to make the *Startup Type* **Automatic**.
- Carbonite Safe Server Backup must be installed and all CSSB operations must be performed as a user with Administrator-level privileges.
- Carbonite Safe Server Backup requires access to TCP Ports 10080 and 10081, which are the default ports used by CSSB for backup and restore operations.
    - Alternate ports can be specified. Please refer to the Administration section in this document.
- CSSB depends on the Visual C++ Redistributable. VC++ must be installed and in working order on the system.
    - Errors that reference *MSVCP120.dll* indicate a damaged or corrupted Microsoft VC++ Redistributable. The systems administrator must repair or reinstall VC++ to resolve the errors.

Supported Backup Types

The following types of data can be backed up and restored by Carbonite Safe Server Backup.

- Windows NTFS files and folders
- Windows ReFS files and folders
- Microsoft SQL Server 2000, 2005, 2008, 2012, 2014 and 2016
- Exchange Server 2003, 2007, 2010, 2013 and 2016 Databases
- Exchange Server 2010, 2013 and 2016 Mailboxes

- Exchange Online hosted for Office 365 Mailboxes
- Microsoft SharePoint Server 2007/WSS 3, 2010 and 2013
    - Standalone and single-server SharePoint farm configurations are supported.
    - Multi-server SharePoint farm configurations are not supported.
- Windows System State (including the registry, certificate server, and active directory information)
- MySQL Server 5.x
- Oracle Server 11i and 11g
- Hyper-V (Windows 2008, 2008R2, 2012, and 2016)
    - Server *core* installations are not supported for any version of Windows.
- MailStore Archives
- Bare Metal Image

SQL Server, Exchange Server, SharePoint, and Hyper-V server configurations are automatically discovered by CSSB.

Getting Started

Quick Start

The *Getting Started* section of the Carbonite Safe Server Backup Knowledgebase contains instructions for quickly setting up and starting Carbonite Safe Server Backup. Click here to navigate to the *Getting Started* page.

Installing CSSB and the Cloud Certificate

Installing Carbonite Safe Server Backup involves two steps. First: installing the software onto your computer. Second: importing the Cloud Certificate, to gain the ability to back up your data to our servers.

Installation

Carbonite Safe Server Backup uses the *InstallShield* installer and self-extracting .exe files for installation. You can download and install Carbonite Safe Server Backup from your Carbonite Server account. Log into your account on our website and click the **Install Carbonite** button (*on this server*) and follow the on-screen instructions. Once the installation is complete, you will have to download and import the cloud certificate before you can start backing up data. For more information about installing Carbonite Safe Server Backup on your computer, please review our Knowledge Base article on installing Carbonite Safe Server Backup.

The amandabackup / CarboniteUser User

A new Windows user labeled *amandabackup / CarboniteUser* is created during the installation of Carbonite Safe Server Backup. CSSB uses *amandabackup / CarboniteUser* for all backup-related tasks.

For more information about the *amandabackup / CarboniteUser* user, please review our Knowledge Base article on the amandabackup / CarboniteUser user.

Importing the Cloud Certificate

CSSB uses a *cloud certificate* to verify your subscription status. A valid subscription is required to perform backups.

Once CSSB is installed, download the cloud certificate from within your account. After you download the certificate, you can easily import it via the CSSB user interface.

1. Log into your account as an administrator at https://account.carbonite.com/.
2. Within the Dashboard, click the **Download backup certificate** button.
3. After you have successfully downloaded the cloud certificate to your computer, click **Cloud > Import Cloud Certificate...** from the menu in the Carbonite Safe Server Backup interface.
4. Locate the cloud certificate on your computer. Once your cloud certificate has been successfully uploaded, you will be ready to perform backup and restore operations to the cloud.

Please refer to our Knowledge Base article on Importing the Cloud Certificate for more detailed information.

Command-Line Installation

CSSB can be installed via the command line. Please refer to our Knowledge Base article on Command Line Installation and Configuration of CSSB.

Terminology

The following terms are used throughout this document.

Backup Types

Carbonite Safe Server Backup supports the backup of many types of data, such as File System, System State, Microsoft SQL Server, and more. These are referred to as backup types. Please refer to the **Supported Backup Types** section for details.

Backup Sets

What to back up, when to back it up, how long to keep it, and other such configuration parameters are collectively referred to as a **backup set**. Each backup set has its own unique name, which is chosen by the user, and may only contain one backup type.

For example, you cannot back up an Exchange database and Windows System state in the same backup set. Instead, you must create separate backup sets for each type of backup.

Working with Backup Sets

Backup sets are displayed in the pane on the left edge of the CSSB user interface. While viewing the *Backup* tab, you can create, edit, activate, deactivate, validate, delete, or clone backup sets. You can also start a backup right away with the *Backup Now* option.

All these options are shown in a toolbar at the top of each backup set. The same operations are also available in the *File* menu or by right-clicking on a backup set icon in the left edge of the CSSB user interface.

Backup Now: Use this option to start a backup right away. This backup will be performed in addition to any regularly scheduled backups.

Enable/Disable a backup set: When a backup set is enabled, its scheduled backups will be performed, and it will be eligible for restore operations. Likewise, scheduled backups will not be performed for Disabled backup sets; however, Disabled backup sets are still eligible for restore operations. *Backup Now* cannot be used for Disabled backup sets. Enabling a Disabled backup set will cause it to back up according to its schedule. Backup sets will be in the Enabled state by default.

A backup set must be saved after it is Enabled or Disabled for the change to take effect.

Validate a backup set: This is a simple way to check whether the essential configuration for the backup set is correct or not. A failed validation indicates that a backup cannot be performed, in which case an error will be returned to the user. Backup set configurations are also automatically validated when the backup set is saved and before any backup run is performed.

Delete a backup set: When the *Delete Backup Set* operation is chosen, CSSB will present the user with the following two options.

- **Delete this backup set configuration and all associated backup data from cloud and disk**
    - Selecting this operation will remove the backup set from CSSB. All data associated with the backup set, on both the local disk and on the cloud, will be deleted.
    - Deletion of data, especially from the cloud, can take a long time. Please be patient.
    - A popup will appear to confirm the choice.
- **Just delete the configuration for this backup set from this computer, but retain the associated backup data**

- Selecting this operation will remove the backup set from CSSB, but all data will be left intact on both local disk and the cloud.
- A popup will appear to confirm the choice.
- The backup set can be recovered by running an Import Existing Backup Set operation.

Backup data that is deleted cannot be recovered.

Copy a backup set: This option will create an exact copy of the current backup set. Schedules can be carried over to the cloned backup set at the user's discretion.

The copied backup set can then be modified as necessary. Users who wish to create multiple, similar backup sets can use the clone feature to reduce setup time.

Backup Levels

Full backups, Differential backups, and Incremental backups are the three **backup levels**. Not all backup types will support all three backup levels.

Full back up: This type of backup will back up all the data associated with the backup set.

Differential backup: This type of backup will back up only the data which changed since the last successful full backup.

Incremental backup: This type of backup will back up only the data which changed since the last successful full, differential, or incremental backup.

For more information on different backup levels, refer to What are Incremental and Differential Backups.

Backup Runs

Every time a manual or scheduled backup is performed, a unique entry is created on the *Report* page. Each of these entries is called a **backup run**. Details are displayed for each backup run, including start time, end time, and backup size. If successful, a backup run will also have an associated archive of backup data.

In most cases, backup runs can be restored, managed, and deleted independently of other backup runs.

Other operations can also have runs. Restore runs, upload runs, and download runs are also possible.

Configure Backup Sets

Every backup set for every backup type is configured with the same basic steps. Users select what to back up, where to put the data, when to back up, and how long to keep the backups.

Unless noted, settings for one backup set do not affect other backup sets.

Add New Backup Set

Click the **add new backup set** button to create a new backup set at any time. Backup sets can also be added via the *File* menu. If backup sets haven't already been created, a message to create a backup set will appear in the middle of the window with an additional button.

Default Settings

Every backup set comes with default settings that represent the best practice for the average user. These default settings may vary from one backup type to another.

The defaults are the best choice in most situations.

What Would You Like to Back up?

When a new backup set is created, a list of available backup types will appear. Choose the type of data you wish to back up.

Not all systems will have all backup types enabled. File System, System State, and MySQL backups are always enabled on all systems. All other backup types will be enabled only if CSSB detects that the services corresponding to each are installed on the system. Backup types will be disabled if CSSB is unable to detect a corresponding service.

Select a backup type to proceed to the next step.

Backing Up

Once a backup type is chosen, the *What would you like to back up?* section is replaced by a section named *Backing up*. Most backup types allow users to select individual files, folders or databases. If possible for the selected backup type, a folder/file tree will appear and allow you to make granular backup selections. Within the tree, place checkmarks next to the items you wish to back up.

For some backup types (like System State), this type of granular selection is not possible. A list of items to be backed up will be displayed for these backup types instead of a folder/file tree.

A dropdown box above the tree shows the type of data to be backed up. Users may choose a different type of data from this menu until the backup set is saved. Once the backup set is saved, the backup type is locked and cannot be changed.

Estimate Backup Size

This button can be found below the tree or list in the *Backing Up* section. Pressing it tells CSSB to estimate the size of a full, uncompressed backup. This estimate does not take various file exclusions into account.

Compression is enabled by default. The actual size of a backup will almost always be smaller than the estimated size due to compression or file exclusions.

Set Your Backup Location

This section controls where backups are stored. It contains three options.

Backup To

Each backup run will save a backup archive to one or more locations, as determined by the selections made in the *Backup To* section. The three options are:

- **Backup to Disk and Cloud:** The backup archive will be saved to disk, then uploaded to the cloud immediately and automatically when the backup is finished. *Backup to Disk and Cloud* is the default choice for all Backup Types and is recommended for all users. This hybrid backup approach allows fast recovery from local disk while safeguarding against disaster by storing the backups in the cloud.
- **Backup to Disk:** The backup archive will be saved to the chosen disk location only. An upload can be scheduled at a later time, if necessary.
- **Backup to Cloud:** The backup archive will be saved to the chosen cloud storage location only. Please refer to our Knowledge Base article on Limitations of the Backup to Cloud Operation for more detailed information.

Disk Storage

Choose the location on disk where you wish to store your backups. This location can be a local disk, an external disk, or a network location.

By default, *amandabackup / CarboniteUser* user must have access to the selected folder. If the backup involves network storage, additional configuration may be required to grant access to *amandabackup / CarboniteUser*. A different user account can be specified in the *Tools; Network Location* menu. Please refer to our Knowledge Base article on How to Use External or Mapped Network Drives with CSSB for more detailed information.

CSSB also supports the use of most NAS devices for storage of local backups. Please refer to our Knowledge Base article on How to Use a NAS Device with CSSB for more detailed information.

Cloud Location

CSSB has data centers in the United States, Canada and Europe. You can choose the location where your data is stored.

Users may change their Cloud Location until the backup set is saved. Once the backup set is saved, the Cloud Location is locked and cannot be changed.

Edit Your Backup Schedule

This section controls when CSSB performs a backup. The default settings are based on our average users' needs.

Every backup level (full, differential, or incremental) requires its own schedule.

Schedule Templates

Alternative schedule templates are available in the *Schedule Templates* dropdown menu. These allow users to quickly and easily switch between different schedule plans.

Each backup type has its own default schedules and schedule templates.

Add Custom Schedules

Custom schedules can be created if defaults do not exist. Select **Add Custom Schedule** from the *Schedule Templates* dropdown menu to add a new schedule. All existing schedules in the current backup set will be removed.

Click the **Add Schedule** button to add a new schedule. Keep in mind that a separate schedule is required for each backup level.

Schedule an Upload

The default *Backup to Disk and Cloud* option will upload backups to the cloud immediately after they complete. If it is necessary to upload later, use the Backup to Disk option, then create a custom schedule. An *Upload* checkbox will appear in the *Edit Schedule* dialogue window. Place a check in the box, select a time, and save the schedule. All backups on disk for the current backup set that have not been uploaded will be uploaded at the time specified.

Edit Existing Schedules

Any backup schedule can be edited regardless of whether it is part of a template or a custom schedule. Double-click the schedule to bring up the *edit schedule* dialogue window.

Delete a Schedule

Click the X icon to the right of any schedule to delete that schedule.

Retention Settings

This section controls how long backups are kept on the disk or in the cloud. Expired backups are purged within 24 hours of expiration.

Backups are only removed by the user defined retention schedule or when manually deleted by a user. Every backup is added to the disk or the cloud without overwriting the previous backups. It is thus important to set a retention policy that suits your needs while also staying within your chosen storage quota.

CSSB allows you to specify a separate retention policy for backups stored on disk and backups stored in the cloud.

While configuring a retention policy, consider the following:

- Importance of older backup data, to ensure that you have access to older files even if those files have been deleted or modified and are not contained in newer backup runs. Past backups also allow for recovery from viruses such as Cryptolocker, which encrypt users' data with unbreakable encryption.
- Free space on disk and/or cloud, to ensure that your backups do not fail due to lack of space.

Retention Types

Users may choose between two types of retention.

Number of Cycles: This option allows users to keep a chosen number of backup cycles. A backup cycle is a full backup plus all associated incremental and differential backups. When a new full backup is completed, excess

backup cycles are purged. Please refer to our Knowledge Base article on Number of Successful Full Backups to Retain for more detailed information.

Time Retained: This option allows users to keep backups for a specific period. However, a backup will not be purged if another backup depends on it. In other words, no backup in a cycle will be purged until *all* the backups in that cycle are ready to be purged. This option ensures that you have all the data that is needed to perform a complete restore to any available backup point, but it does mean that backups can be kept longer than their retention setting suggests.

Special Retention Values

CSSB offers two special retention values.

Forever: Change the retention value to **0** (zero) to keep backups forever. A backup set to be kept forever will never expire. It will never be purged automatically, but it can be deleted manually. Remember that backups do not overwrite each other. Please ensure that you have ample space on disk and/or cloud if you choose to keep backups forever.

Delete After Upload: Only available for disk backup retention, this causes the backup on disk to be deleted immediately after it is uploaded to the cloud. Only backups that have successfully uploaded will be deleted. Set your retention value to **-1** to use the *Delete After Upload* option.

Advanced Backup Settings

This section controls advanced settings.

Data Encryption

CSSB offers two different types of encryption.

Auto Encryption

Carbonite Safe Server Backup will automatically encrypt your backups on the cloud using AES-128 bit (or better) encryption.

Transfer of your backup data to the cloud is secure using the Transport Layer Security (TLS) protocol.

The encryption itself takes place server-side on the cloud. Local backups will not be encrypted if Auto Encryption is chosen. If encryption of locally-stored backups is required, you must use the *Private Key Encryption* option.

Private Key Encryption

With Private Key Encryption, users create a unique key generated from a passphrase. The encryption key or passphrase is required to restore data.

Private Encryption Keys will encrypt your backups using AES-256 bit encryption. AES-256 bit encryption is trusted worldwide.

If Private Key Encryption is chosen, you are responsible for safe and secure storage of your encryption keys. CSSB does not store your encryption keys or passphrase anywhere. Replacement keys can be created if you remember your passphrase. Replacement keys can be used to decrypt backups.

If you lose your private encryption key and forget your passphrase, neither you nor Carbonite will be able to decrypt your encrypted backup data.

Important Information about Private Key Encryption

There is no way to transfer a key from one system to another for the purposes of continuing backups. Nor can you import an existing key to continue backups.

In all such situations, you must create a new key. If you use the same passphrase, the new key will be identical to the old key. This avoids potential problems where some backups in a backup cycle are encrypted with one key, and some with another.

You *may* use an existing key to decrypt backups during a restore.

Create a New Private Encryption Key

First, create a personal encryption key. To create the key:

1. Click the **Edit** button to the right of the *Advanced backup settings* panel.
2. Click **Add Private Key with 256-bit encryption**.
3. A new window will appear.
4. Enter a Key Name. This will be the name of the file stored on disk.
5. Choose a passphrase.
    a. The passphrase must be at least four characters in length.
6. Confirm the passphrase.
7. Choose a location to which your encryption key will be saved.
    a. The key must remain in this location to encrypt your backups.
8. Click **OK**. A new window will appear to confirm that the encryption key was created.
9. Click **OK** again in the new window.
10. Save the backup set. All future backups for this backup set, local and cloud, will be encrypted using your Private Encryption Key.
    a. Past backups, if they exist, are not retroactively encrypted.

Add an Existing Private Encryption Key to a Backup Set

You may use an already-existing Private Encryption Key with other backup sets.

1. Click the **Edit** button to the right of the *Advanced backup settings* panel.
2. Click **Add Private Key with 256-bit encryption**.
3. A new window will appear and ask you if you wish to use the existing key.
4. Click **Enable**.
5. Save the backup set. All future backups for this backup set, local and cloud, will be encrypted using your Private Encryption Key.
    a. Past backups, if they exist, are not retroactively encrypted.

Delete an Existing Private Encryption Key

To delete an encryption key:

1. Select any backup set for which Private Key Encryption is enabled.
2. Click the **Edit** button to the right of the *Advanced backup settings* panel.
3. Click **Delete Private Key**.
4. A window will appear to confirm your deletion and warn you that a copy of the key should be kept for decryption purposes.

Existing Private Encryption Keys cannot be modified. To change a key, first delete a key and create a new one.

Compression

Compression makes the backups smaller. A small backup can be uploaded to the cloud much faster than a large one. This can help save network bandwidth and reduce the storage space required for backups.

Compression is ON by default, but it can be turned OFF.

CSSB cannot accurately predict how much any given backup set will compress. Some data, such as text documents and databases, will compress very well. Other data, like most multimedia files, does not compress well.

Bandwidth Throttling

By default, CSSB attempts to transfer data to and from the cloud at the maximum possible speed. In other words, CSSB will use all the bandwidth it can.

Bandwidth throttling allows users to limit the amount of bandwidth used by CSSB. Throttling can be set at a global level (affecting all backup sets) or configured for a specific backup set.

Users may select a maximum bandwidth allotment separately for upload and for download. Also, users may choose to have throttling always apply, or to only apply at certain times.

All throttle speeds are in kilobits per second. A value of 0 means the speed is unlimited.

Throttle by Speed

The *Throttle by Speed* option sets a maximum bandwidth allotment that is always on.

Throttle by Time

This option allows you to throttle only at certain times and on certain days. For example, it is possible to throttle from 8am to 5pm every weekday. Time-based throttling is highly configurable; values can be set in 15-minute increments.

Click the buttons for Max Upload Rates or Max Download Rates to begin configuration. A time table will appear. Select the cells from the time table shown and input the desired value in the *Maximum Rate (kbps)* box.
Click **Apply Custom Rate to Selected Slots** to apply the provided value to the selected slots. Please note that configured values take effect only for future data transfers. Any ongoing jobs will continue to use any previously configured values.

For ease of use, CSSB comes with predefined templates for time-based bandwidth throttling. You can choose a template from the *Choose Predefined Template* dropdown. If you wish to modify the template values, you can select the time slots as described above and configure different values.

Custom Scripts

CSSB can run batch scripts before and/or after any given backup. There are strict requirements for batch scripts used in this manner. Please refer to our Knowledge Base article on How to Use Batch Scripts with CSSB for more detailed information.

Email Notification Preferences

An email can be sent to report the status of backups, uploads, downloads, and restores for all backup types, with one exception.

Email notifications will not be sent upon success or failure of any restore of the Bare Metal Image backup type. Emails can be sent for all other operations in the Bare Metal Image backup type.

A SMTP server must be specified. Click the **Edit** button to the right of the *Email notification preferences* panel or navigate to **Preferences > Email** to specify SMTP server information. Once configured, email notifications are enabled for all backup sets.

By default, emails are sent only when an operation has failed. Emails can also be sent after a success or a warning.

Backup Details and Requirements

Every backup type has its own requirements. Application and database backups must meet the proper requirements. Some backup types also have special configuration options.

Back Up the Windows File System

Windows File System refers to files and folders stored on local, external, or network disks.

Overview

File System backups have the following properties:

- NTFS and ReFS file systems are supported.
- There are no restrictions on file size.
- Any file of any extension can be backed up. This includes virtual hard disks, such as those used for virtual machines.
- Open, locked, and in-use files can be backed up if files are on a local disk (including external hard drives) and the Volume Shadowcopy Service (VSS) is started.
- Backup of mapped network drives is supported.
    - Open, locked, and in-use files will not be backed up.
    - Please refer to our Knowledge Base article on How to Use External or Mapped Network Drives with CSSB for more information.
- Backup of NAS devices is supported.
    - Open, locked, and in-use files will not be backed up.
    - Please refer to our Knowledge Base article on How to Use a NAS Device with CSSB for more information.
- CSSB can be used to backup and restore folders that participate in DFS replication topology.
    - CSSB must be installed on the host where the replicated DFS share is configured. CSSB relies on the DFSR VSS Writer for correct functionality of backup and restore operation.
    - Backup and restore of a network shared folder which is a member of a DFS replication group is not supported.
- Junction Point Directories in Windows Vista, Server 2008, or higher will not be backed up to ensure that backups do not include the same data multiple times.
    - Please ensure that the source directory is included in your backup set.
    - Please refer to our Knowledge Base article on Files Under Junction Point Directories are not Backed Up for more information.

- CSSB will automatically exclude all files and folders contained in its own directories from File System Backups.
- CSSB will also exclude files and folders that are dedicated to CSSB's use, even if they are not in the default installation path. Examples include:
    - Folders where local backups are stored will be excluded, so that CSSB will not back up its own backups recursively.
    - Folders chosen as a Restore location, such as the global location chosen in the *Advanced* menu, so that performing a restore does not cause backup size to balloon.
- Certain system files are automatically excluded from File System backups. These files are included as part of a System State backup.
    - Please refer to our Knowledge Base article on File Types Excluded from File System Backup for more information.

Backup Levels

File System backups support Full, Differential, and Incremental backups.

Special Options

The following special options are available for File System backups.

Exclude Files

Files and folders can be excluded from a File System backup. Exclusions are listed below the file tree.

Click the dropdown menu to see pre-configured exclusion lists, or manually enter a list of files or folders.

Wildcards can be used with exclusions. Wildcards include:

- The asterisk (*), which replaces any number of characters.
- The question mark (?), which replaces a single character.

Exclusions should be separated with a space.

Example

Here is an example list of exclusions: ***.pst C:\Folder* *$Recycle* file?.txt C:\Folder\document.doc**.

- The ***.pst** entry excludes any file with the .pst extension.
- The **C:\Folder*** entry excludes all files and folders (including subfolders) under C:\Folder\.
- The ***$Recycle*** entry excludes any file or folder that contains $Recycle in the path.

- The **file?.txt** entry excludes any file that has any single character between "file" and ".txt", such as file1.txt or filex.txt. File32.txt would not be excluded.
- The **C:\Folder\document.doc** entry only excludes the single file at C:\Folder\document.doc.

CSSB and Deduplication Volumes in Windows Server 2012 and higher

In CSSB 4.14 and lower

CSSB cannot properly back up deduplication volumes in version 4.14 and lower. Although the backup will complete without an error message, files will not be backed up. The total size of the backup will be far smaller than expected.

In CSSB 5.0

Attempts to back up files and folders on a deduplication volume will fail with an error that deduplication volumes are not supported. You must deselect any files on a deduplication volume and save the backup set.

In CSSB 5.1 and higher

We support Windows native deduplication in Windows Server 2012 and higher.

Requirements

The following requirements must be met for Windows File System backups to function.

- The Volume Shadowcopy Service must be started and functional for open, locked, or in-use files to be backed up.
- The *amandabackup / CarboniteUser* user must have access to the files to be backed up.
    - If the files are in a network location, you may instead specify a different user account in *Tools; Network Location*. The account you choose must have access to all the files to be backed up.

Requirements for VSS Snapshots on a SMB3 Network Share

CSSB can take snapshots of files located on network shares configured with the SMB3 protocol. This allows backup of open, in-use, and locked files that exist on these shares.

The following requirements must be met.

- The following backup types allow VSS snapshots of data on a SMB3 network share:
    - File System
    - Microsoft SQL Server
    - Hyper-V (Full backups only)
- The application server on which Carbonite Safe Server Backup is installed and the file server that houses the data must both be running Windows Server 2012 or higher.

- The application server and file server must be joined to the same Active Directory domain.
- The File Server VSS Agent Service role service must be installed on the file server.
- The *amandabackup* / *CarboniteUser* user (or the user specified in **Tools; Network Location**) must be a member of the Backup Operators and Administrators groups on both the application server and file server.
    - The *amandabackup* / *CarboniteUser* user is granted these roles by default during CSSB installation. If a different user account is specified in **Tools; Network Location**, please ensure this user has the correct roles.
        - If *amandabackup* / *CarboniteUser* (or the user specified in **Tools; Network Location**) has been created as a domain user, simply ensure that *amandabackup* / *CarboniteUser* is a member of both groups on both machines.
            - In some environments, you may have to add *amandabackup* / *CarboniteUser* (or the user specified in **Tools; Network Location**) to the local Backup Operators and Administrators groups instead of using the domain groups.
        - If *amandabackup* / *CarboniteUser* (or the user specified in **Tools; Network Location**) has been created as a local user on the CSSB system and does not exist on both machines, you must manually create an identical user account on any system where it does not exist.
            - Use the same password for all systems.
            - Ensure that the user account is a member of the Backup Operators and Administrators groups on all systems.
- The *amandabackup* / *CarboniteUser* user (or the user specified in **Tools; Network Location**) must have read-only or greater access to the file server.
    - It is recommended to give Full Control to this user account.

A snapshot cannot be created if these requirements are not met. If you are attempting a Hyper-V or Microsoft SQL Server backup, the backup will fail immediately without a snapshot.

If you are attempting a File System backup, CSSB will attempt to copy the files even if a snapshot cannot be created. Open, locked, and in-use files will fail to back up, but the backup will progress unless too many files cannot be backed up. Please refer to this Knowledge Base article.

Back Up the Windows System State

Windows System State refers to a collection of several key operating system elements and their files. Backing up the Windows System State is crucial for a successful disaster recovery strategy.

Overview

Windows System State backups have the following properties:

- System State backups will contain the items in the list below (if present on the system). Components that are not installed will not be backed up. For example, if Active Directory is not installed, then Active Directory will not be included in the System State backup:
    - **Boot Files:**
        - For Windows versions older than Vista: **SystemDrive\NTDETECT.COM**, **SystemDrive\ntldr**, **SystemDrive\boot.ini** (**System Drive** is usually **C:**).
        - For Windows Vista and newer: **SystemRoot\boot** directory (**SystemRoot** is usually **C:\Windows**)
    - **Catalog Files: SystemRoot\System32\CatRoot.**
    - **MachineKeys Files: SystemRoot\System32\Microsoft\Protect*** and **AllUsersProfile\ApplicationData\Microsoft\Crypto\RSA\MachineKeys***, where **AllUsersProfile** is **C:\Documents and Settings\All Users**.
    - **Performance Counters**: perf*.dat and perf*.bak files in **C:\Windows\System32** on all Windows versions.
    - **WFP Files**: All .dll and .exe files that come under Windows File Protection (WFP). Usually the .dll files reside in **C:\windows\system32**.
    - **IIS Metadata File**: If IIS is installed (applicable to all Windows versions).
    - **Certificate Database** (Applicable to only Windows 2003 server that are Certificate Servers): Files in **C:\Windows\System32\certsrv**.
    - **COM+** registration database.
    - **Registry**: System, default, SAM, Security and Software files in **SystemRoot\System32\config** and additional Components and Schema files in Vista.
    - **Active Directory**: If Active Directory is installed, backups include the database, log files, and Group Policy Objects (GPOs).
- System State components cannot be chosen individually. System State will always back up all components currently installed on the system.
- Files included in a System State backup are automatically excluded from File System backups. Please refer to our Knowledge Base article on File Types Excluded from File System Backup for more information.

Backup Levels
Windows System State backups are always a full backup. Differential and incremental backups are not supported.

Requirements
The following requirements must be met for Windows System State backups to function.

- System State backups and Microsoft SQL Server backups should not be run simultaneously.
- The Volume Shadowcopy Service must be started.
- System state backups of a Windows Domain controller that holds the Active Directory Certificate Services role may encounter an "Application's Writer is Not Available" error.
    - Please refer to this Microsoft Knowledge Base article to resolve the problem.
- All local volumes attached to the system must be of the NTFS or ReFS file system. If a local volume is of any other file system, the backup will fail.

Back Up a Microsoft SQL Server
A *Microsoft SQL Server* is a database application. It can stand alone or be part of another application.

Overview
Microsoft SQL Server backups have the following properties:

- Microsoft SQL Servers are automatically discovered at backup set creation.
- Individual databases can be selected for backup.
- Only databases are backed up.
    - Other MSSQL files (such as installation files) are not backed up. These files must be included in a File System backup set to back them up.
- Transaction logs are truncated during any full or incremental (log) backup. Transaction logs are not truncated during differential backups.
- Only databases that are in the Mounted state will be backed up.
- System databases are only backed up during full backups. They are skipped during differential and incremental (log) backups.
- Read-only databases are only backed up during full backups. They are skipped during differential and incremental (log) backups.
- Only local databases can be backed up. Backups of databases on another system are not supported.
- New databases will be automatically detected and backed up if all databases are selected for backup.

- Databases that are removed from a Microsoft SQL Server will be automatically removed from the backup if all databases are selected for backup.
- Databases are not automatically added or removed from the backup if individual databases are selected.

Recovery Models

Databases configured with the SIMPLE recovery model, FULL recovery model, and BULK recovery model are supported.

Simple Recovery Model

- The full backup will contain .MDF, .LDF and .NDF (in case of filegroups) files in the backup image.
- The differential backup will contain .LDF files.
- Incremental (Log) backups will be skipped for databases using the Simple Recovery Model.

Full Recovery Model

- The full backup will contain .MDF, .LDF and .NDF (in case of filegroups) files in the backup image.
- The differential backup will contain the changed blocks of the .MDF database file.
- Incremental (Log) backups will contain .TRN files (transaction logs flushed to the disk). These are transactions that have changed since the last backup of any level, be it Full, Differential, or Incremental.

Bulk Recovery Model

- The full backup will contain .MDF, .LDF and .NDF (in case of filegroups) files in the backup image.
- The differential backup will contain the changed blocks of the .MDF database file.
- Incremental (Log) backups will contain .TRN files (transaction logs flushed to the disk). These are transactions that have changed since the last backup of any level, be it Full, Differential, or Incremental.

Backup Levels

Microsoft SQL Server backups support Full, Differential, and Incremental (log) backups.

Special Options

The following special options are available for Microsoft SQL Server backups.

Granting Access to Microsoft SQL Server

CSSB will attempt to automatically grant access to your Microsoft SQL Server instances when you create a Microsoft SQL Server backup set. This access is granted to the *Carboniteuser* or *amandabackup* service account, whichever is applicable in your environment. Access must be granted to all Microsoft SQL Server instances selected for backup. Backups will fail if access is not granted.

If access cannot be granted automatically, an error will appear stating that access was denied. The error will include a link to this Knowledge Base article, and it will also state which of your Microsoft SQL Server instances have denied access.

Follow the steps in the article to grant access.

Checking Access to Microsoft SQL Server

After you save your Microsoft SQL Server backup set, a *Check Access* button will appear below the list of databases on the left side of the screen.

Click this button if you wish to verify that proper access has been granted. For example, you can click this button after granting permissions manually.

An error will appear if access is denied to any of your Microsoft SQL Server instances.

Requirements

The following requirements must be met for Microsoft SQL Server backups to function. Unmet requirements are the most common cause of MSSQL backup failures.

- The *amandabackup / CarboniteUser* user must be added as a SQL server user with enough privileges. Please refer to our Knowledge Base article on Allowing access to Microsoft SQL Server for more information.
- The Volume Shadowcopy Service must be started.
- Microsoft SQL Server and Windows System State backups should not be performed simultaneously.
- The SQL VSS Writer Service must be running at the time of backup. CSSB will attempt to start the service if it is disabled.
- TCP/IP must be enabled for all MSSQL instances to be backed up. TCP/IP settings are controlled in the SQL Server Configuration Manager tool.
- Transaction log-based databases such as Microsoft SQL Server, Exchange, or SharePoint are not intended to be backed up by multiple backup applications.
    - Incremental and differential backups will usually fail if multiple applications are used on the same database.
    - In some cases, incremental and differential backups may not fail outright, but still aren't able to be restored.
    - Please refer to our Knowledge Base article on Databases and Multiple Backup Applications for more information.

CSSB can perform differential backups of SQL Server 2014 databases with memory-optimized tables. However, preview builds of SQL 2014 are not optimized for differential backups.

In CTP2 (Community Technology Preview 2) or earlier, the differential backup will include all data and delta files as though it were a full backup. There will be no reduction in size; a differential is essentially the same thing as a full back up in this circumstance.

SQL 2014 RTM (release-to-manufacturing) is optimized so that differential backups are smaller than full backups in the expected manner.

Please see this Microsoft Blog for more information: Differential Database Backup with Memory-Optimized Tables

Microsoft SharePoint databases can be backed up as if they were Microsoft SQL Server databases. If possible, use the SharePoint backup type to back up SharePoint databases.

Do not back up Microsoft SharePoint databases in more than one backup set. The result is the same as if multiple backup applications were used. Please refer to our Knowledge Base article on Databases and Multiple Backup Applications for more information.

Requirements for VSS Snapshots on a SMB3 Network Share

CSSB can take snapshots of files located on network shares configured with the SMB3 protocol. This allows backup of open, in-use, and locked files that exist on these shares.

The following requirements must be met.

- The following backup types allow VSS snapshots of data on a SMB3 network share:
    - File System
    - Microsoft SQL Server
    - Hyper-V (Full backups only)
- The application server on which Carbonite Safe Server Backup is installed and the file server that houses the data must both be running Windows Server 2012 or higher.
- The application server and file server must be joined to the same Active Directory domain.
- The File Server VSS Agent Service role service must be installed on the file server.
- The *amandabackup* / *CarboniteUser* user (or the user specified in **Tools; Network Location**) must be a member of the Backup Operators and Administrators groups on both the application server and file server.

        - The *amandabackup* / *CarboniteUser* user is granted these roles by default during CSSB installation. If a different user account is specified in **Tools; Network Location**, please ensure this user has the correct roles.
            - If *amandabackup* / *CarboniteUser* (or the user specified in **Tools; Network Location**) has been created as a domain user, simply ensure that *amandabackup* / *CarboniteUser* is a member of both groups on both machines.
                - In some environments, you may have to add *amandabackup* / *CarboniteUser* (or the user specified in **Tools; Network Location**) to the local Backup Operators and Administrators groups instead of using the domain groups.
            - If *amandabackup* / *CarboniteUser* (or the user specified in **Tools; Network Location**) has been created as a local user on the CSSB system and does not exist on both machines, you must manually create an identical user account on any system where it does not exist.
                - Use the same password for all systems.
                - Ensure that the user account is a member of the Backup Operators and Administrators groups on all systems.
- The *amandabackup* / *CarboniteUser* user (or the user specified in **Tools; Network Location**) must have read-only or greater access to the file server.
    - It is recommended to give Full Control to this user account.

A snapshot cannot be created if these requirements are not met. If you are attempting a Hyper-V or Microsoft SQL Server backup, the backup will fail immediately without a snapshot.

If you are attempting a File System backup, CSSB will attempt to copy the files even if a snapshot cannot be created. Open, locked, and in-use files will fail to back up, but the backup will progress unless too many files cannot be backed up. Please refer to this Knowledge Base article.

Back Up a Microsoft Exchange Server Database

A *Microsoft Exchange Server* is an email-based communications server.

Overview

Microsoft Exchange Server backups have the following properties:

- VSS based backups of Store database files (.edb & .stm), transaction logs, and checkpoint files for all the Mounted Storage Groups are performed.
- Mailbox Databases and Public Folder Databases can be backed up.
- New databases will be automatically detected and backed up if all databases are selected for backup.

- Databases that are removed from Exchange will be automatically removed from the backup if all databases are selected for backup.
- Databases are not automatically added or removed from the backup if individual databases are selected.

Backup Levels

Microsoft Exchange Server backups support Full, Differential, and Incremental backups.

- **Full Backups** contain the Mailbox database (.edb & .stm files), transaction logs, and checkpoint files.
    - After a successful Full Backup, Exchange will purge all committed transaction logs.
    - An event is logged by Exchange in the Windows Application Event Logs when transaction logs are purged.
- **Incremental Backups** contain the transaction logs only.
    - After a successful Incremental backup, Exchange will purge all committed transaction logs.
    - An event is logged by Exchange in the Windows Application Event Logs when transaction logs are purged.
- **Differential Backups** contain the transaction logs only.
    - No logs are truncated or purged after a successful Differential backup.

Requirements

The following requirements must be met for Microsoft Exchange Server backups to function.

- Backups for all versions of Exchange except Exchange 2003 have additional dependencies.
    - Powershell 2.0 or higher must be installed on the system.
    - .NET Framework 4.0 or higher must be installed.
    - Use the **Check Dependencies** option in the *Tools* menu to verify that these are installed.
- The Volume Shadow Copy service must be started.
- In Windows Small Business Server 2003, the Exchange VSS Writer is *disabled* by default.
    - CSSB will attempt to enable the VSS writer automatically. However, the Microsoft Exchange Information Store service must be restarted manually before backup will work.
    - The writer can be enabled manually by following these instructions in this Microsoft knowledgebase: How to turn on the Exchange writer for the Volume Shadow Copy service in Windows Small Business Server 2003
- The latest Service Packs for Microsoft Exchange must be installed.

- Only Mounted Storage Groups are backed up. Unmounted Storage Groups are not backed up.
- The amandabackup / CarboniteUser user must be assigned to the proper role in the Exchange Security Groups in Exchange 2007 and higher. CSSB will attempt to assign the proper role automatically.
    - These roles are:
        - *For Exchange 2007*: View Only Organization Management
        - *For Exchange 2010 and up*: Exchange View-Only Administrators
    - Please follow these steps if the *amandabackup / CarboniteUser* user must be added to the Exchange Security Groups manually.
        - Open "Active Directory Users and Computers"
            - For Exchange 2007: Add the *amandabackup / CarboniteUser* user to *Microsoft Exchange Security Groups > View Only Organization Management*
            - For Exchange 2010 and 2013: Add the *amandabackup / CarboniteUser* user to *Microsoft Exchange Security Groups > Exchange View-Only Administrators*
        - The Carbonite Safe Server Backup Services must be restarted after *amandabackup / CarboniteUser* is added to the Exchange Security Groups
            - Click on **Tools > Restart Services** in the CSSB User Interface, or restart *Carbonite Safe Server Cloud Controller* and *Carbonite Safe Server Backup Controller* manually from the Windows Services Panel.
- Circular logging must be disabled within Exchange. Circular logging can interfere with backup operations and is not recommended in a normal production environment:
    - In most cases, you *will not* need to disable Circular logging manually because CSSB will do it automatically.
    - If it is necessary to disable Circular Logging manually, please refer to the Microsoft articles below:
        - Exchange Server 2007: http://technet.microsoft.com/en-us/library/bb331968%28v=exchg.80%29.aspx
        - Exchange Server 2010: http://technet.microsoft.com/en-us/library/dd297937.aspx
- Transaction log-based databases such as Microsoft SQL Server, Exchange, or SharePoint are not intended to be backed up by multiple backup applications.
    - Incremental and differential backups will usually fail if multiple applications are used on the same database.

- In some cases, incremental and differential backups may not fail outright but still aren't able to be restored.
- Please refer to our Knowledge Base article on Databases and Multiple Backup Applications for more information.

Back Up a Microsoft SharePoint Server

A *Microsoft SharePoint Server* is a web application framework that integrates intranet, content management, and document management.

Overview

Microsoft SharePoint Server backups have the following properties:

- Standalone and single-server farms are supported. Multi-server farms are not supported.
- The following SharePoint components contained in the SQL database and application web site files are backed up (if present on the system). The components listed below cannot be selected or deselected for backup individually:
    - Configuration and Admin databases
    - Content and configuration data for Web Applications (Eg: C:\Program Files\Microsoft Office Servers\14.0\Data\Office Server\Applications)
    - Any third-party databases that are registered with SharePoint
    - Shared services databases in SharePoint
    - Office Search & Help Search index files
- Other objects, such as Site collection, Web site, List/Document library, Document library folder, Document library file, List item, and Version, are not backed up by a SharePoint backup set.
    - Files such as the SharePoint installation directory, IIS metabase information, Website application pool directory, and so on must be protected by other backup sets.
        - Files, documents, the folders that contain them, and the SharePoint installation directory itself should be backed up as part of a File System backup.
        - IIS metabase information, port configurations, host header, authentication, and other such configuration data is included as part of a System State backup.
            - Please refer to our Knowledge Base article on Applications for more information.

Backup Levels

Microsoft SharePoint Server backups support Full and Differential backups. Incremental (log) backups are not supported.

Requirements

The following requirements must be met for Microsoft SharePoint Server backups to function.

- The Volume Shadow Copy Service must be started.
- The SQL VSS Writer service must be running at the time of backup and recovery. CSSB will attempt to automatically start the SQL VSS Writer service.
- The Windows SharePoint Services VSS Writer service will be registered and started automatically by CSSB.
- The SharePoint VSS writer depends on other VSS writers, such as OSearch and SPSearch. These writers must be enabled and set to automatic start for SharePoint backups to complete successfully.
    - For SharePoint 2010
        - SPSearch4 VSS Writer
            - Enable by starting the SharePoint Foundation Search V4 service.
        - OSearch14 VSS Writer
            - Enable by starting the SharePoint Server Search 14 service.
    - For SharePoint 2007
        - OSearch VSS Writer
            - Enable by starting the Office SharePoint Server Search service.
        - SPSearch VSS Writer
            - Enable by starting the Windows SharePoint Services Search service.
- Windows SharePoint Services VSS Writer service must run under the admin app pool account, which is the Network Service account in a basic installation of Windows SharePoint Services.
- Transaction log-based databases such as Microsoft SQL Server, Exchange, or SharePoint are not intended to be backed up by multiple backup applications.
    - Incremental and differential backups will usually fail if multiple applications are used on the same database.
    - In some cases, incremental and differential backups will not fail, but cannot be restored.

- Please refer to our Knowledge Base article on Databases and Multiple Backup Applications for more information.

Special Requirement for WSS 3.0

The default installation of Windows SharePoint Services 3 instances will use the Windows Internal Database (MICROSOFT##SSEE), which is run under the "Network Service" system account.

In these cases, the "Network Service" account must have Full Control permissions to the "misc" directory found in the folder where you installed CSSB.

By default, the misc directory is located here:

- For Windows Server 2003: C:\Program Files\Carbonite\Carbonite Safe Server Backup\misc\
- For Windows Server 2008 and 2012: C:\ProgramData\Carbonite\Carbonite Safe Server Backup(x64)\misc\

Microsoft SharePoint databases can be backed up as if they were Microsoft SQL Server databases. If possible, use the SharePoint backup type to back up SharePoint databases.

Do not back up Microsoft SharePoint databases in more than one backup set. The result is the same as if multiple backup applications were used. Please refer to our Knowledge Base article on Databases and Multiple Backup Applications for more information.

Back Up a MySQL Server

MySQL Server is a database application.

Overview

MySQL Server backups have the following properties:

- Logical backups are performed using the mysqldump utility.
- Individual databases within a MySQL server can be selected for backup.
- If a single database is selected, specific tables within that database can be backed up.
    - Create multiple MySQL backup sets if specific tables from multiple databases should be backed up.
- Backups can be taken of local MySQL servers or remote MySQL servers located on another machine.
- MySQL Servers, unlike other databases, are not automatically discovered by CSSB. As a result, the MySQL Server backup type will always be an available option when creating a new backup set.

Backup Levels

MySQL Server backups are always Full backups. Differential and Incremental backups are not supported.

Special Options

The following special options are available for MySQL Server backups.

Discovery

MySQL Servers are not automatically discovered. Information about the server must be provided for a connection to be established and a backup taken.

- *MySQL User:* Enter the username of a MySQL user with enough access to take the MySQL Server backup.
- *Password:* Enter the password for the MySQL User.
- *MySQL Server:* Enter the IP address of your MySQL Server. A "localhost" value refers to the local machine.
- *Port:* Enter the port on which MySQL is listening.
- *"mysqldump" parameters:* Additional parameters can be passed to the mysqldump utility. Enter them here if they are needed.
- **Path to mysql.exe:* Browse to the location where mysql.exe and mysqldump.exe are located.

The following defaults are assumed:

- *MySQL Server:* localhost
- *Port:* 3306
- *"mysqldump" Parameters:* --lock-all-tables

These defaults work for most MySQL instances. Two common exceptions are listed below:

- If the MySQL instance uses a socket file for connection instead of a port, please leave the Port field blank.
- MySQL instances that use the INNODB storage engine may find that replacing the *--lock-all-tables* parameter with the *--single-transaction* parameter can result in increased performance and lower CPU usage.

Once configured, click the Discover button to connect to the MySQL server. A list of databases and tables will appear.

Estimate Size

The Estimate Size button is not available for MySQL backups.

Requirements

The following requirements must be met for MySQL Server backups to function.

- The MySQL user specified in the Discovery section must have enough privileges to perform backup and restore. The minimum set of privileges are:
    - For backup:
        - LOCK TABLES, SELECT, FILE, RELOAD, SUPER, UPDATE, TRIGGER, SHOW VIEW
    - For restore:
        - CREATE, DROP, INDEX, SHUTDOWN, INSERT, ALTER, UPDATE, TRIGGER, SUPER, REPLICATION CLIENT, CREATE VIEW
- MySQL client utilities (*mysqldump* and *mysql*) must be installed on the CSSB machine, and the MySQL client version must be compatible with MySQL server.

Back Up an Oracle Server

An *Oracle Server* is a database application sold and supported by Oracle.

Overview

Oracle Server backups have the following properties:

- The Oracle database, Control File, Server Parameter file, all table spaces, and Archived logs are backed up.
- All backup files & archived logs from the Flash Recovery Area are also included in the backup.
- The online redo logs will not be included in the backup.
- Databases, log files, control files, etc. cannot be selected for backup individually.

Backup Levels

Oracle Server backups support Full, Differential, and Incremental backups.

Requirements

The following requirements must be met for Oracle Server backups to function.

- The Oracle VSS Writer service must be installed for the Oracle instances that need to be backed up.
    - The Oracle VSS Writer is provided by Oracle and is not part of a default Windows server.
    - If the Oracle VSS Writer service is not installed, please refer to the Oracle documentation.
- The Oracle VSS Writer service must be started.
- The Oracle VSS Writer service must be set to Automatic start.
- The Volume Shadow Copy Service must be started.
- All Oracle databases in *NOARCHIVELOG* mode must be *Mounted* and in *Read-Only* state.

        - The database cannot be open in *Read-Write* mode. If it is, the backup will fail.
        - A backup may contain multiple databases. If one database is in *Read-Only* state and the others are in *Read-Write* state, the backup for all databases will fail.
- Oracle databases in *ARCHIVELOG* mode may be open in Read-Write or in *Mounted* state.

Back Up Hyper-V

Hyper-V is Microsoft's virtualization host. It allows users to create, run, and manage virtual machines.

Overview

Hyper-V backups have the following properties:

- Windows "Core" Servers running Hyper-V are not supported.
- Individual Virtual Machines may be selected for backup.
    - All virtual hard disk files associated with the virtual machine will be displayed. However, you cannot choose individual VHD files. If you select one VHD within a virtual machine, all VHD files will be automatically selected.
- The Initial Store (InitialStore.xml) is always backed up and will not appear in the "Backing Up" list.
- VMs are backed up by one of two methods.
    - The "Child VM Snapshot" method uses the Volume Shadowcopy Service (VSS) inside a child VM. This allows the VM to be backed up without forcing it into a saved state.
        - This is the default backup method. If a Child VM Snapshot cannot be taken, a Saved State backup is taken instead.
        - Child VM Shapshot backups have additional requirements, as outlined below.
    - The "Saved State" method briefly places each VM into a saved state so that a snapshot can be taken. Once the snapshot is taken, the VM is returned to its previous state.
        - This method is used if the Child VM Snapshot method cannot be performed.
- New VMs will be automatically detected and backed up if all VMs are selected for backup.
- VMs that are removed from Hyper-V will be automatically removed from the backup if all VMs are selected for backup.
- VMs are not automatically added or removed from the backup if individual VMs are selected.
- Virtual hard disk files located on network shares can only be backed up if they are on a SMB3 network share. VHD files on other network shares cannot be backed up.

- If the guest VM is running Microsoft SQL Server, Exchange, or SharePoint, then transaction logs for those databases will be truncated when a Hyper-V backup occurs.
    - Truncation will occur as a part of both Full and Incremental backups.
    - Truncation will occur during both Saved State and Child Snapshot backups.
    - In short, truncation will occur during any Hyper-V backup. This can affect the ability to perform backups of these databases from within the guest VM.
    - Users should take care when performing backups of their MSSQL, Exchange, and SharePoint databases from within the guest VM if they are performing backups from the Hyper-V host.
        - The Hyper-V backup, because it truncates the transaction logs, will cause incremental and/or differential backups of MSSQL, Exchange, and SharePoint to fail as described in Databases and Multiple Backup Applications.

Backup Levels

Hyper-V backups support Full backups on Windows Server 2008 and higher.

Incremental backups are supported on Windows Server 2008 R2 and higher.

Differential backups are not supported on any platform.

Special Options

The following special options are available for Hyper-V backups:

- *Back up a running VM only if its hot backup can be performed* is a special option designed to keep VMs running smoothly during backup.

If enabled, CSSB will skip backup of a running VM if a "Child VM Snapshot" backup cannot be performed.

Requirements

The following requirements must be met for Hyper-V backups to function.

- The Hyper-V VSS Writer must be started.
- Cluster Shared Volumes are not supported.
- Restore of Hyper-V virtual machines and Host Configuration must be done to the *same* version of Windows as took the backup.
    - For example, if the original Hyper-V backup was taken on Windows Server 2008 R2, the restore must be done to Windows Server 2008 R2.
- CSSB will automatically install and start the Carbonite Safe Server Hyper-V Service on 64-bit systems running Hyper-V when CSSB 4.14 or higher is installed.

    - This includes upgrades from previous versions.
    - This service allows CSSB to take and maintain incremental backups for Hyper-V. Any interruption will cause incremental backups to fail, as described below.
        - The Carbonite Safe Server Hyper-V Service must be running at the time of the most recent successful full backup. If it was not running at that time, attempts to perform incremental backups will fail.
        - If the service is stopped, restarted, or suffers any interruption after the full backup, subsequent incremental backups will fail until a new full backup is performed.
        - The errors associated with these failures will prompt the user to run a new full backup.
- If there are multiple Hyper-V backup sets on the system, they can be run simultaneously.
    - However, the backups should be started at least one minute apart to allow time for VSS Snapshot creation.
    - A virtual machine should not be backed up in more than one backup set.

Additional Requirements for Child VM Snapshot backups

The following requirements must be met for Child VM Snapshot backups to function.

- Backup (volume snapshot) Integration Service must be installed and running on the child VM.
    - The service name is *Hyper-V Volume Shadow Copy Requestor*.
- The child VM must be in the running state. VMs that are not running always use the Saved State backup method.
- The Snapshot File Location for the VM must be located on the same volume in the host operating system as the VHD files.
    - Example: If the VHD files are on C:\, the Snapshot File Location must be on C:\
- All volumes in the child VM must be basic disks. Dynamic disks are not supported by the Child VM Snapshot method.
- All disks in the child VM must use a file system that supports snapshots, such as NTFS or ReFS.

Requirements for VSS Snapshots on a SMB3 Network Share

CSSB can take snapshots of files located on network shares configured with the SMB3 protocol. This allows backup of open, in-use, and locked files that exist on these shares.

The following requirements must be met.

- The following backup types allow VSS snapshots of data on a SMB3 network share:
    - File System
    - Microsoft SQL Server
    - Hyper-V (Full backups only)
- The application server on which Carbonite Safe Server Backup is installed and the file server that houses the data must both be running Windows Server 2012 or higher.
- The application server and file server must be joined to the same Active Directory domain.
- The File Server VSS Agent Service role service must be installed on the file server.
- The *amandabackup* / *CarboniteUser* user (or the user specified in **Tools; Network Location**) must be a member of the Backup Operators and Administrators groups on both the application server and file server.
    - The *amandabackup* / *CarboniteUser* user is granted these roles by default during CSSB installation. If a different user account is specified in **Tools; Network Location**, please ensure this user has the correct roles.
        - If *amandabackup* / *CarboniteUser* (or the user specified in **Tools; Network Location**) has been created as a domain user, simply ensure that *amandabackup* / *CarboniteUser* is a member of both groups on both machines.
            - In some environments, you may have to add *amandabackup* / *CarboniteUser* (or the user specified in **Tools; Network Location**) to the local Backup Operators and Administrators groups instead of using the domain groups.
        - If *amandabackup* / *CarboniteUser* (or the user specified in **Tools; Network Location**) has been created as a local user on the CSSB system and does not exist on both machines, you must manually create an identical user account on any system where it does not exist.
            - Use the same password for all systems.
            - Ensure that the user account is a member of the Backup Operators and Administrators groups on all systems.
- The *amandabackup* / *CarboniteUser* user (or the user specified in **Tools; Network Location**) must have read-only or greater access to the file server.
    - It is recommended to give Full Control to this user account.

A snapshot cannot be created if these requirements are not met. If you are attempting a Hyper-V or Microsoft SQL Server backup, the backup will fail immediately without a snapshot.

If you are attempting a File System backup, CSSB will attempt to copy the files even if a snapshot cannot be created. Open, locked, and in-use files will fail to back up, but the backup will progress unless too many files cannot be backed up. Please refer to this Knowledge Base article.

Back Up Exchange Local Mailboxes

An *Exchange Local Mailbox* backup allows users to select one or more mailboxes from a local Exchange server for backup and restore, including an individual mailbox.

An Exchange Local Mailbox backup is intended to supplement an Exchange Database backup; not replace it. Transaction logs are only purged as part of an Exchange Database backup, which is also superior for disaster recovery purposes. Perform both an Exchange Local Mailbox backup and an Exchange Database backup for best results.

Overview

Exchange Local Mailbox backups have the following properties:

- Once user credentials are supplied, CSSB will automatically discover and display all available mailboxes.
- Users may select all mailboxes or individual mailboxes for backup and restore.
- Emails, Calendars, Contacts, Contact Groups, and Tasks will be backed up.
    - Attendees and attachments are not included in Calendar backups.
- Items in the Junk folder will not be backed up.
- Items in Resource mailboxes will not be backed up.
- Items in Equipment mailboxes will not be backed up.
- Items in Public Contact Groups will not be backed up.
- New mailboxes will be automatically added to the next full backup if *all* mailboxes are selected. New Mailboxes will not appear in incremental backups until the full backup is complete.
- Mailboxes that are removed from Exchange will be automatically removed from the backup if *all* mailboxes are selected.
- Mailboxes are not automatically added or removed from the backup if individual mailboxes are selected.
- Transaction logs are *not* purged as part of an Exchange Local Mailbox backup. Add an Exchange Database backup set to manage the Exchange transaction logs.

Backup Levels

Exchange Local Mailbox backups support Full and Incremental backups. Differential backups are not supported.

Requirements

The following requirements must be met for Exchange Local Mailbox backups to function.

- CSSB will attempt to use Exchange Autodiscovery to detect and connect to the Exchange server.
    - Autodiscovery must be configured and accessible from the system for automatic detection of your Exchange server.
        - Microsoft provides a free tool to test Autodiscovery at https://testconnectivity.microsoft.com.
        - CSSB attempts the Autodiscovery URL methods, which are included in the above test.
            - Only the URL test results are relevant. CSSB cannot connect to Exchange if the URL tests fail.
            - CSSB cannot connect to Autodiscover via the “HTTP redirect” or “DNS SRV redirect” methods.
    - If Autodiscovery fails, you will be prompted to enter the URL of your Exchange server.
        - Please refer to this Knowledge Base article on how to find your Exchange Server URL.
    - Either Autodiscovery must work or you must supply the correct Exchange URL. CSSB cannot connect to Exchange without this information.
- Only local Exchange servers are supported. To back up Exchange Online, use the Exchange Online Mailbox backup type instead.
- The system must have Exchange 2010 or Exchange 2013 installed. Exchange 2007 and earlier are not supported.
- The most recent service packs for Exchange and Windows must be installed.
    - For example, Exchange 2010 Service Pack 3 is available. Therefore Exchange 2010 running Service Pack 2 is not supported.
- An administrator account must be specified within CSSB. The username must be in the User Principal Name (UPN) or email format, such as UserName@Domain. For example, Admin@company.com and username@company.onmicrosoft.com are both in the correct format.
    - For all Exchange versions, the Exchange Impersonation right must be granted to this administrator account. CSSB will attempt to add the Impersonation right automatically.

- For Exchange 2013, sometimes the mailboxes will not be discovered and displayed after the account is specified. If this occurs, grant the eDiscovery permission to the administrator account.
    - Please see http://technet.microsoft.com/en-us/library/dd298059(v=exchg.150).aspx

Back Up Exchange Online hosted for Office 365

An *Exchange Online hosted for Office 365* backup allows users to select one or more mailboxes from an Exchange Online server for backup and restore, including an individual mailbox. Exchange Online is part of Microsoft Office 365.

An Exchange Online hosted for Office 365 backup is the only way to back up and restore Exchange Online with CSSB. The Exchange Database and Exchange Local Mailbox backup types cannot access Exchange Online.

Overview

Backups of Exchange Online hosted for Office 365 have the following properties:

- Once user credentials are supplied, CSSB will automatically discover and display all available mailboxes.
- Users may select all mailboxes or individual mailboxes for backup and restore.
- Emails, Calendars, Contacts, Contact Groups, and Tasks will be backed up.
    - Attendees and attachments are not included in Calendar backups.
- Items in the Junk folder will not be backed up.
- Items in Resource mailboxes will not be backed up.
- Items in Equipment mailboxes will not be backed up.
- Items in Public Contact Groups will not be backed up.
- New mailboxes will be automatically added to the next full backup if *all* mailboxes are selected. New Mailboxes will not appear in incremental backups until the full backup is complete.
- Mailboxes that are removed from Exchange will be automatically removed from the backup if *all* mailboxes are selected.
- Mailboxes are not automatically added or removed from the backup if individual mailboxes are selected.

Backup Levels

Exchange Online hosted for Office 365 backups support Full and Incremental backups. Differential backups are not supported.

Requirements

The following requirements must be met for Exchange Online Mailbox backups to be performed:

- Exchange Online Mailbox backups are not supported on Windows Server 2003, Windows XP, or Windows Vista.
- An administrator account must be specified within CSSB. This user must have the Global Administrator role. The username must be in the UPN (User Principal Name) or email format, such as *UserName@Domain*.
    - For example, *Admin@company.com* and *username@company.onmicrosoft.com* are both in the correct format.
- The Exchange Impersonation right must be granted to this administrator account. CSSB will attempt to add the Impersonation right automatically.
    - This right can be granted to an account from the Admin page in the Office 365 portal.
- The Exchange Discovery Management role must be granted to this administrator account.
    - This role can be granted to an account from the Admin page in the Office 365 portal.
- CSSB uses Exchange Autodiscovery to detect and connect to the Exchange server. The Autodiscovery connection attempt is made based on the administrator account provided.
    - If Autodiscovery fails, CSSB cannot connect to Exchange Online.
    - Microsoft provides a free tool to test Autodiscovery at https://testconnectivity.microsoft.com.
    - It is not necessary to pass a specific test. CSSB can connect to Exchange Online using the URL, DNS SRV redirect, and HTTP redirect methods.

Back Up a MailStore Archive

MailStore is a Carbonite-owned company that specializes in email archiving. Carbonite Safe Server Backup can back up and restoring MailStore archives.

Overview

MailStore backups have the following properties:

- MailStore installations are automatically discovered and the MailStore backup type will be available when creating a new backup set.
- MailStore archive components will be listed on the *Backup* page.
    - Individual components cannot be added or removed from backup. For example, one cannot choose to remove the Master Database component from the backup set.

Backup Levels

MailStore backups support Full, Differential, and Incremental backups.

Requirements

- The MailStore VSS Writer must be installed, running, and functioning correctly.
- MailStore archives that are stored in a network location cannot be backed up.

Back Up a Bare Metal Image

A *Bare Metal Image* backup allows for the restore of an entire system to a similar or dissimilar hardware. They are highly useful for migrations and disaster recovery.

Overview

Bare Metal Image backups have the following properties:

- Bare Metal backups are only available with certain Server plans.
    - If a user downgrades from a plan that supports Bare Metal Image backups to a plan without it:
        - Bare Metal Image backups will fail with an error that *Bare Metal Image backups are not supported by your subscription.*
        - The backup images on disk and/or cloud will *not* be automatically deleted.
- Only entire volumes can be selected for backup, but it is possible to restore individual files and folders.
- Exchange, Microsoft SQL Server, and SharePoint databases are not included in Bare Metal Image backups.
    - A separate backup set will be automatically created for each of these applications, if that application exists on a drive that is selected for Bare Metal Image backup.
        - Please refer to the *Bare Metal Suite* section of this Knowledge Base article for more information.
    - All other data types, such as MySQL, MailStore, and Hyper-V are included in the Bare Metal Image backup set.
- Only one Bare Metal Image backup set can be created per system.
    - If another Bare Metal Image backup set is imported using the instructions in Importing Existing Backup Sets, and a Bare Metal backup set already exists, you must choose which backup set remains enabled.
        - You will be asked which backup set you wish to enable.
        - All other Bare Metal Image backup sets will be disabled automatically.
- New volumes added to the system will not be automatically selected for backup.
- Disks that contain LVM and LDM partitions are not supported and cannot be backed up or restored.

- Attempts to back up or restore these disks will fail. The error will specify that LVM and LDM disks cannot be backed up or restored.

Backup Levels

Bare Metal Image backups support Full and Incremental backups. Differential backups are not supported.

Note: The size of the Incremental backups is directly related to the number of files within the backup. The more files there are, the larger the incremental backups.

Special Options

Bare Metal Image backup sets behave quite differently from other backup sets. There are several new options and a few restrictions that apply only to Bare Metal Image backup sets.

- The option to backup to cloud only, without saving to disk first, is not available for Bare Metal Image backups.
    - You can still upload the backups to cloud with an *Upload from Disk* schedule, or by using the *Backup to Disk and Cloud* option, just like any other backup type.
- The disk storage location must meet certain requirements, as noted below.
- A *Recovery Media* is required to perform a complete Bare Metal Image restore.
    - You will be prompted to create your Recovery Media when you save your Bare Metal Image backup set for the first time.
- A Bare Metal Image backup is always compressed. All compression options are disabled.
- Bare Metal Image backups on disk cannot be encrypted at this time, even if Private Key Encryption is used.
    - Both automatic and Private Key Encryption will work for backups that are uploaded to the cloud.

Requirements

The following requirements must be met for Bare Metal Image backups to function:

- BMI backups and restores cannot be performed on any 32-bit hardware, regardless of the version of Windows. However, you can take BMI backups on a Virtual Machine that is emulating 32-bit hardware, so long as the physical VM host machine has 64-bit hardware.
    - The Windows Operating System itself can be 32-bit, unless it is Windows XP or Windows Server 2003.
        - For Example:
            - A 32-bit Windows Server 2008 system with 32-bit hardware cannot perform a BMI backup or restore, because the hardware is 32-bit.

            - However, a 32-bit Windows Server 2008 system installed on 64-bit hardware will be able to perform BMI backups and restores. The bit-level of Windows does not matter for Server 2008.
    - On Windows XP and Server 2003, BMI backups will not function if the Windows Operating System itself is 32-bit.
        - For Example:
            - A 32-bit Windows Server 2003 system on a 64-bit hardware cannot take BMI backups, even though the hardware is 64-bit.
- The chosen disk storage location must meet these requirements.
    - The path to the disk storage location must not be longer than 45 characters.
    - The disk storage location cannot be on a volume you have chosen to back up as part of a Bare Metal Image backup.
        - For example, if you are backing up the C:\ drive, you cannot choose a disk storage location anywhere on the C:\ drive.
        - Network storage locations and external disks are recommended storage for Bare Metal Image backups.
- Additional software is automatically installed when a Bare Metal Image backup set is saved for the first time. This process usually only takes a few minutes.
    - If the installation of this additional software is canceled, interrupted, or fails, Bare Metal Image backups will not function.
- A minimum of 2GB of memory is required for backup and restore for Bare Metal Image backups.
    - However, this requirement increases according to the number of files in the backup and the number of incremental backups between full backups.
        - Please refer to the *Best Practices* section of this Knowledge Base article for more information on memory and performance of Bare Metal Image backups.
- Windows must be running the English or German language pack. Core functionality, such as creating Recovery Media may fail if any other language is used.
- Bare Metal Image backups cannot back up volumes whose volume name include parentheses.
    - For example, if your volume name is “OS (Windows Server 2012)”, backups will fail.
    - The error message may include one or both error codes:

i.ZWC_INVALID_BACKUP_COMPONENT

ii.ZWC_MSSQL_BACKUP_LIST_CONTAIN_INVALID_DB

- To back up these volumes, you must remove the parentheses from your volume names.

- BMI backups cannot function on a system that has enabled the Federal Information Processing Standard (FIPS) setting.
    - This *only* applies to BMI backups. All other backup types work on FIPS-enabled systems.
    - FIPS is a compliance setting targeted at government facilities as described here: https://technet.microsoft.com/en-us/library/cc180745.aspx.
    - While FIPS is enabled, the additional software required for BMI backups cannot be installed.
        - An event will be logged. It calls out FIPS as the reason.
            - Service cannot be started. System.InvalidOperationException: This implementation is not part of the Windows Platform FIPS validated cryptographic algorithms.
    - Do not temporarily disable FIPS to install the Bare Metal Image software.
        - Bare Metal Backup and Restore of FIPS-enabled system is not supported.
        - If FIPS cannot be permanently disabled, you must back up use other backup types.
- Volumes with full-disk encryption cannot be backed up in a Bare Metal Image backup set.
    - Bitlocker encrypted volumes will not appear in the list of volumes to be selected.
    - Volumes with any other type of full-disk encryption, such as DiskCryptor, will be visible and can be selected. However, all attempts to back up these disks will fail.
- You cannot select network locations for BMI backup.

Note: There is no fixed size requirement for the recovery media target device. In many cases, a USB thumb drive as small as 512 MB may be enough. However, the exact size depends upon the options chosen during the creation of the recovery media. For example, adding device drivers increases the size of the recovery media. The recovery media creation wizard will inform you if the chosen target device is too small.

For additional information about performing a Bare Metal Image backup, backup best practices, and creating a recovery media, please refer to this Knowledge Base article.

Restore Details and Requirements

There are some considerations to take when restoring from Carbonite Safe Server Backup.

Effects of Retention on Restore

Any backup that is set to be purged by retention will still be purged, even if that backup is currently involved in a restore. Users who stick to the default retention of 2 (or more) backup cycles are unlikely to encounter this scenario. Users who set their retention are far more likely to encounter it.

To avoid purging backups mid-restore, you can edit your backup schedule in the backup set settings or update the retention of the relevant backups in the History section.

Restore Requirements

The following requirements must be met for all restores regardless of whether the target is the original machine or an alternate machine:

- The *amandabackup / CarboniteUser* user must have full access to the folder where the backups are stored.
- For cloud backups, this is the chosen Download Folder.
- For local backups, this is the folder where backups are stored on the disk.
- If the backups are stored in a network location, *amandabackup / CarboniteUser* (or the user account specified in *Tools; Network Location*) must have access to the network folder. Please refer to our Knowledge Base article on How to Use External or Mapped Network Drives with CSSB for more information.
- Likewise, the *amandabackup / CarboniteUser* user must have full access to the folder(s) to which the backups are being restored.

Additional Requirements for All Restores to an Alternate Machine

Backups can be restored from one machine to another. Restores to an alternate machine are most frequently performed for testing purposes or as a part of disaster recovery scenario where the original machine is no longer accessible.

When restoring to another machine, the following steps must be performed before the restore can begin.

- Install CSSB on the new machine.
- Import the cloud certificate.
- Import Existing Backup Sets

Additional Requirements for Application Restores

Restores to an alternate machine of applications or databases, such as Hyper-V or Microsoft SQL Server, frequently have additional requirements that must be met. Please refer to the restore section that corresponds to the application or database for more information about application restores.

Special *metadata* is stored for every backup run. The metadata is known as the **backup catalog** and includes information about when a backup was taken, what is inside the backup, what backup level it is, and much more.

When restoring to an alternate machine, the backup catalogs must be restored before the data can be restored.

The **Import Existing Backup Sets** operation scans for backup runs available on disk or cloud and restores the backup catalog. This rebuilds the backup set(s) to which the backups belong and includes the backup set configuration, the Restore page, and the Report history.

You can Import Existing Backup Sets from backups on disk or cloud.

Import Existing Backup Sets

1. Install CSSB and import the cloud certificate.
2. Click **Tools > Import Existing Backup Sets**.
3. Enter information about the backups you wish to recover. Only backups that match the criteria will be imported.
4. Click **OK**.
5. A scan for valid, matching backups will begin.
    - This process can take a long time. Please be patient.
6. When prompted, close and restart CSSB to complete the Import Existing Backup Sets process.
7. All functions, including reports and data restores, will be available for the backups and backup sets.

Backup sets can only be imported if the backup set does not already exist on the system. If it is necessary to Import Existing Backup Sets for a backup set that already exists, please delete the existing backup set before you begin.

Import Existing Backup Sets from Cloud

When Importing Existing Backup Sets from the cloud, CSSB must search your cloud backups and download the backup set information. The actual backups themselves are not downloaded. Only backup sets that match the criteria below will be restored.

Cloud Location

Choose the cloud datacenter where the backups are stored. The default is *All Cloud Locations*. Restricting this setting to a specific cloud location can greatly speed up the process.

Host Name

This refers to the host name of the original machine that took the backups. Do not include the Windows Domain name, if any.

- Leave this field blank to import backup sets from the current machine.
- Enter an asterisk (*) to import backup sets for all machines on the account.
- Enter a specific host name to import backup sets for that machine only.

Backup Set Name

This refers to the name of the Backup Set that needs to be restored.

- Leave this field blank to import all backup sets.
- Enter a backup set name to import that backup set only.
- An asterisk (*) can be used as a wildcard if only part of the backup set name is known.
    - For example, entering **exchange** into this field would import any backup set that contains the word "exchange."

Restore Location

Backup set information must be downloaded to the local machine before it can be imported. Specify a directory here to use as temporary storage of this data. This data is deleted from disk at the end of the Import Existing Backup Sets process.

Import Existing Backup Sets from Local Directory

You can restore the backup sets from backup files kept on a local disk. Only backup sets that match the criteria below will be imported.

Host Name

This refers to the host name of the original machine that took the backups. Do not include the Windows Domain name, if any.

- Leave this field blank to import backup sets from the current machine.
- Enter an asterisk (*) to import backup sets for all machines on the account.
- Enter a specific host name to import backup sets for that machine only.

Backup Set Name

This refers to the name of the Backup Set that needs to be imported.

- Leave this field blank to import all backup sets.
- Enter a backup set name to import that backup set only.
- An asterisk (*) can be used as a wildcard if only part of the backup set name is known.
    - For example, entering **exchange** into this field would import any backup set that contains the word "exchange."

Backup Data Location

This is the directory in which the local backups are stored. You must specify the exact folder in which the local backups are stored. Sub-folders will not be scanned for backups. If backups are in multiple directories, perform multiple Import Existing Backup Set operations.

Every restore is configured with the same basic steps. Users choose a point in time, which data to restore, where the data is restored, and what to do if there is a conflict.

Backups on disk will always be used for if they are available and accessible. The restore process will automatically download the backups from the cloud if the disk backups are not available.

Choose The restore Point

A list of all available backup runs for the current backup set is shown. Select one for restore. This is the **restore point**.

The restore point can be further refined in the dropdown box below the list. This dropdown box displays all the previous backup runs in the same cycle as the chosen restore point.

Choose **All Latest** to restore the latest version of data, as of the selected Restore Point, that is contained within any full, differential, or incremental backup in the backup cycle. This option eliminates the need to perform multiple restores to get the latest data. It is not necessary to restore from the full backup, and then the first incremental, and then the second incremental, and so on. When *All Latest* is chosen, the latest data from all backup runs is restored. All Latest is the default setting for restores.

If an individual backup is chosen instead of All Latest, only the data from the selected backup run will be restored. Data that is in other backup runs will not be restored.

Select Data to Restore

Data can be selected for restore in three ways:

Restore All

The *Restore All* option selects all data at the chosen Restore Point for restore. Restore All is the default selection.

Restore Select

The *Restore Select* option allows users to select some data for restore instead of all data. The Restore Select button will only appear if it is supported by the current backup type.

A file tree will be displayed for File System backups. A list or summary will be displayed for databases, applications, and for any backup of any type that was taken before Carbonite Safe Server Backup version 4.12.

Place checkmarks next to the items that should be restored. Items that are selected will be restored. Items that are not selected will not be restored.

Search

The *Search* option allows users to search for individual files, folders, or databases. The Search button will only appear if it is supported by the type of backup you have selected for restore.

You may select multiple items in the Search window, but the Search option is best when selecting a small number of items for restore. If you need to restore many items, consider using *Restore Select* instead.

Versions

If multiple versions of an item are backed up, a plus (+) sign will appear next to the item. Click the plus sign to expand the list of versions.

By default, *all versions* of the matching items will be shown from all backup runs. You can filter your version results from the *Show versions from* dropdown at the top of the Search window.

Download Size During Restore Select and Search

It is uncommon for CSSB to require an entire backup to be downloaded for a selective restore. The sole exception is for Bare Metal Image backups, which *always* require the entire backup to be present on disk before restoring can begin, even for selective restores. Total restores of all data in a backup require the entire archive, of course.

All CSSB backups are packaged into an archive file, much like a zip file. When performing a Selective Restore from a backup on the cloud, CSSB does not normally need to download the entire backup archive(s) to be restored. Instead, it can download certain pieces of the archive (known as *chunks*) and extract the data from those. This prevents the need to download multiple gigabytes or even terabytes of data just to restore one file.

Each chunk is a maximum of 10MB in size. At restore time, CSSB will grab only those chunks it needs to restore. Typically, this includes some number of chunks necessary to open the archive, such as headers and metadata plus the chunks of the archive that contain the data which was selected for restore. Selective restore of multiple files often increases the number of chunks required.

The necessary chunks are downloaded and automatically merged into a file. Then CSSB can open and extract the data from this file.

This process is automatic and largely invisible to the user. However, it's important to note that it is not uncommon for a selective restore to download tens or even hundreds of megabytes more than the size of the data selected for restore. Final download size depends on how many chunks CSSB needs to download to create a complete archive file on disk from which it can extract the data.

Review The restore Settings

Additional restore settings can be modified by clicking the **Edit** button.

Restore To

Most backup types allow restore to the *original location* or to an *alternate location*. Some application or database restores will have different options that replace the standard *Restore To* options.

Restore Folder

The *Restore Folder* is the location where your data will be restored. You *must* choose a folder the first time you perform a restore for any backup set. Future restores will default to the same folder used by the previous restore.

You may change the folder from the *Restore* page at any time.

Download Folder

When there are no backups on disk, or if the backups on disk cannot be accessed, the backup archive(s) will be downloaded from the cloud and stored in this folder. The disk(s) must have enough space to store the backup archive(s) and the data selected for restore.

Name Conflict Settings

This option controls what happens if a file already exists while a restore is in progress. There are four possible choices.

- **Overwrite Existing:** The existing file will be overwritten by the file in the backup.
- **Keep Existing:** The existing file will be kept. The file from the backup will not be restored.
- **Rename Existing:** The existing file will be renamed by appending a timestamp to its file name. The restored file will be restored with its name intact.
- **Rename Restored:** The restored file will be renamed by appending a timestamp to its file name. The existing file will keep its name intact.

Keep Downloaded Archive

Backup archives that are downloaded for restore are automatically deleted from disk when the restore is complete. Enable this option to keep the backup archive on disk.

This option is primarily used to reduce download time if multiple restore operations must be run from the same backup.

The backup archive is stored in the Download Folder and must be deleted manually from disk to reclaim space.

Perform Archive Verification

When enabled, the integrity of the backup archive will be scanned to ensure the archive is not corrupted or otherwise damaged.

Archive verification can take a very long time. You should only select this option if verifying the data integrity is more important to you than restoring all possible files as soon as possible.

Run Script Before/After Restore

Scripts can be run before or after any restore. Please refer to our Knowledge Base article on How to Use Batch Scripts with CSSB for more information.

Encryption Options

If Private Key Encryption was used, special options will appear on the *Restore* page. Please refer to How to use encryption for more details on how to restore encrypted backups. Please refer to our Knowledge Base article on Encryption for more information.

Restore Now

Click the Restore Now button to begin the restore. You will be prompted to confirm your choices. The Monitor page will appear once the restore begins.

Restoring the Windows File System

Restores of the *Windows File System* are among the easiest and most flexible offered by CSSB. Files can be selected individually or in groups and restored to local disks, network drives, and alternate machines with few restrictions and no additional requirements.

Functionality

CSSB will restore the chosen files to the chosen location.

Additional Requirements

There are no additional requirements.

Restoring the Windows System State

The *Windows System State* is designed primarily for disaster recovery by restoring the registry, boot files, system files, Active Directory, and more. When used correctly in combination with a File System backup, a machine can be brought back to the exact state it was in at the time of backup.

An example walkthrough of a complete system recovery, including File System and System State, can be found in our Knowledge Base article on Recovery of a Windows 7 Computer.

Functionality

The entire contents of the Windows System State backup will be restored. Selective restore of individual System State components is not possible.

By default, System State will restore to the Original Location and Overwrite Existing files.

Additional Requirements

System State restores have additional requirements.

- The Windows version, including Service Packs, of the restore machine must match the original machine where the System State backup was taken.

        - For example, if the original machine was Windows Server 2003 with Service Pack 2, the machine accepting the restore must also be Windows Server 2003 with Service Pack 2.
        - Any variation will result in failure to restore.
- It is strongly recommended to restore the System State to identical hardware as the original machine.
        - Restore to similar hardware is usually possible but may result in errors or failure to restore.
        - The less similar the hardware to the original, the less likely it is that the System State restore will succeed.

System State restores to a Domain Controller with the Active Directory role can be very complex. Please refer to How to Restore System State on an Active Directory Domain Controller for full details.

Restoring a Microsoft SQL Server

Microsoft SQL Server backups can be restored to the original machine or to an alternate machine.

Additional Requirements

Microsoft SQL Server restores have many additional requirements, including some requirements unique to specific Restore Locations.

Requirements for All Microsoft SQL Server Restores

- Microsoft SQL Server must be installed on the system.
- Microsoft SQL Server must be started and healthy unless the *rebuild system databases* option is checked. Please see below for more details.
- Do not run backup and restore operations of an MS-SQL backup simultaneously.
- SQL 2005 restores of system databases (such as model, master, and msdb) will fail if other applications are actively connected to the MS-SQL server. Please disable any SQL query analyzers, the SQL Management Studio, and other similar programs before performing a restore of system databases.
- The SQL Server VSS Writer service must be running at the time of backup and recovery. Microsoft recommends that the SQL VSS Writer service be automatically started. MSDE writer is not enough for backup and recovery.
- The *amandabackup / CarboniteUser* user must have access to the SQL server. There are two ways to grant access:
        - Click the **Allow Access** button on the Backup page of any Microsoft SQL backup set
        - Add the user in SQL Server Management Studio:
            - Click **Security > Logins > Add New Login**. Add the *amandabackup / CarboniteUser* user account, then add it to the *sysadmin server* role.

- The log on user for the SQL Server service must have full permissions to the folder that was chosen for restore.
    - Locate this information in *Services.msc*. The user will be listed as the log on user for the SQL Server service that matches the instance(s) that are to be restored.
- System databases (such as model, master, and msdb) should be restored separately before restoring user databases.

Requirements for Restoring to an Alternate Machine

The *Restore a Copy of Database to Original or New Location* and *Restore to a New Location and Overwrite Original Database* options, discussed below, may both be used to restore to an alternate machine:

- The SQL Server instance name must be the same on the new machine as it was on the original machine. The restore will fail if the instance names do not match.
- Restore of system databases require that the SQL Server version exactly match between the new machine and the original machine. The restore will fail if there is a version mismatch.
    - This includes any restore that contains the "master", "model", and/or "msdb" databases, including restores in which the *Rebuild System Databases* option is enabled.
    - Example: The model database of an SQL Server 2005 database cannot be restored to anything other than a SQL 2005 instance.
- Restore of user databases require that the SQL Server version on the new machine be the same as or more recent than the original machine.
    - Example: A user database from an SQL 2008 instance can be restored to another SQL 2008 instance or a SQL 2012 instance. It cannot be restored to a SQL 2005 instance.
- Restore of user databases require that the target machine the same (or more recent) version of the Windows operating system as the original.
    - Example: A SQL 2008 instance running on Windows Server 2003 can be restored to a SQL 2008 instance running on Windows Server 2003, 2008, or 2012.
    - Example: A SQL 2008 instance running on Windows Server 2008 cannot be restored to Windows Server 2003. It can be restored to Windows Server 2012.

System databases such as model, master, and msdb, should *only* be restored to the original machine, an identical server, or a new server following a bare metal restore using CSSB's Bare Metal Image backup type.

The Microsoft SQL Server may not start if system databases are restored to an entirely different server, because these databases contain configuration data concerning MSSQL itself. This includes restores where users select the *rebuild system databases* option.

Restoration of system databases to a non-identical server requires significant database administration knowledge and is outside of the scope of Carbonite Safe Server Backup.

Special Options

There are several special options to consider when restoring a Microsoft SQL Server.

Run DBCC CHECKDB After Restore

When selected, this option verifies the integrity of the database(s) after restoring. It is off by default.

Rebuild System Databases

This option is intended to be used following a bare metal restore. It may also be used in other situations where Microsoft SQL Server is installed but cannot be started.

The Microsoft SQL Server cannot be started after a bare metal restore because the system databases are not included in the bare metal recovery process. CSSB cannot restore system or user databases unless MSSQL is running.

Place a check mark next to *Rebuild System Databases* to rebuild the system databases during restore. This option will replace the existing system databases with the system databases from the backup. Once complete, CSSB will attempt to start the MSSQL server.

Once MSSQL starts, CSSB will restore the other databases selected for restore. Any databases that were selected for restore will be restored after the rebuild, automatically. There is no need to run a separate restore.

Please see this Knowledge Base article for more information.

Restore To

The *Restore To* options have been replaced for Microsoft SQL Server restores with the following options:

- Original Location
- Restore a Copy of Database to Original or New Location
- Restore to a New Location and Overwrite Original Database

Each Restore Location may have additional requirements beyond those found above. Each Restore Location option is designed to fill a specific purpose.

Restore To: Original Location

This option will restore the selected database to the location where it was originally located at the time of backup. There are no additional requirements.

If the database to be restored is currently attached to the SQL instance, and the current location of the database files is *different* from the time of backup, then CSSB will treat the current database location as the "original" location.

Restore To: Restore a Copy of Database to Original or New Location

This method allows a user to restore the selected database with a new name to either its original location or to a completely new location. This is like creating a copy of a database with a new name in a new location.

Information

- SQL recovery is performed during the restore.
- The database file names (.ldf & .mdf) will remain the same as they were before.

Additional Requirements

- System databases such as master, model, and msdb, cannot be restored using this method.
- To restore the database to the Original location, the Path field must be left blank.
- When restoring to the Original location using this method, the Original database must be deleted or detached from SQL before the restore begins.
    - If the Original database still exists, the restore will fail.
- Once a user selects on the Restore page, he cannot make any changes to the selection list unless he toggles between Restore Methods.

Procedure to Restore a Copy of a Database to its Original Location

In the following example, we describe a database named "Sales" that needs to be renamed to "SalesTeam" and restored to its original location.

- Go to the *Restore* page and select "Sales" database in the File Path view.
    - It is not necessary to select the database in all the full, differential, and/or incremental backup runs. If one selects the "Sales" database in any one backup run, CSSB will automatically select the "Sales" database from all other backup runs that are required for the restore operation.
- From the *Restore To* drop down, select **Restore a Copy of Database to Original or New Location** option.
- The *Edit Database Name & Restore Path* dialog box will appear.
    - Specify the new name ("SalesTeam") in *New Name* field.
    - Keep the *Path* field blank. This tells CSSB to restore the database to its original location.
    - Save the changes.
- (Optional) Select the *Run DBCC CHECKDB* check box if you wish to verify the logical and physical integrity of all objects within the specified database(s) after the restore completes.
- Click the **Restore** button.

- A confirmation box will appear. Click **OK** to restore or **Cancel** to cancel.
- After the restore completes, the "SalesTeam" database will appear in the SQL Management Studio.

Procedure to Restore a Copy of a Database to a New Location

In the following example, we describe how to make a copy of a database named "Marketing" that originally exists at *C:\Program Files\Microsoft SQL Server\MSSQL10.MSSQLSERVER\MSSQL\DATA* with a name "MarketingTeam" that will be created at *E:\Microsoft SQL Server\MSSQL10.MSSQLSERVER\MSSQL\DATA*.

- Go to the *Restore* page and select the "Marketing" database in the File Path view.
    - It is not necessary to select the database in all the full, differential, and/or incremental backup runs. If one selects the "Marketing" database in any one backup run, CSSB will automatically select the "Marketing" database from all other backup runs that are required for restore.
- From the *Restore To* drop down menu, select **Restore a Copy of Database to Original or New Location** option.
- The *Edit Database Name & Restore Path* dialog box will appear.
    - Input the new name ("MarketingTeam") in the *New Name* field.
    - Click the *Path* field and choose a restore folder.
    - Save the changes.
- (Optional) Select the *Run DBCC CHECKDB* check box if you wish to verify the logical and physical integrity of all the objects in the specified database(s) after the restore completes.
- Click on the **Restore** button.
- A confirmation box will appear. Click **OK** to begin the restore or **Cancel** to cancel.
- After the restore completes, both the "Marketing" & "MarketingTeam' databases will appear in the SQL Management Studio.

Restore To: Restore to a New Location and Overwrite Original Database

This method allows a user to move the selected database to a completely new location. The original database will be overwritten as part of the restore process.

This option may still be chosen if the original database is not present. The Restore to a New Location and Overwrite Original Database option will behave exactly like Restore a Copy of Database to Original or New Location in such a situation.

Information

- SQL recovery is performed during the restore.

- The database file names (.ldf & .mdf) remain the same as before.
- CSSB restores the database to the path specified by the user in the *Edit Restore Path* dialog box.
- Details about the restore procedure are logged by the SQL Server in the Windows Event Viewer.

Additional Requirements

- System databases like master, model, and msdb cannot be restored using this method.
- If the selected database already exists on the SQL server before the restore procedure is initiated, the database files from the original location will be deleted and restored to the new location.
- Once a user selects on the Restore page, he cannot make any changes to the selection list unless he toggles between Restore Methods.

Procedure to Restore to a New Location and Overwrite Original Database

In the following example, we describe a database called "Finance" that originally existed at *C:\Program Files\Microsoft SQL Server\MSSQL10.MSSQLSERVER\MSSQL\DATA*.

This database needs to be moved to a new location at *E:\Microsoft SQL Server\MSSQL10.MSSQLSERVER\MSSQL\DATA*.

- Go to the *Restore* page and select the "Finance" database from within the File Path view.
    - It is not necessary to select the database in all the full, differential, and/or incremental backup runs. If one selects the "Finance" database in any one backup run, CSSB will automatically select "Finance" database from all other backup runs required for the restore.
- From the *Restore To* drop down menu, select the **Restore to a New Location and Overwrite Original Database** option.
- The *Edit Restore Path* dialog box will appear.
    - Click inside the *Path* field and choose a restore folder.
    - Save the changes.
- (Optional) Select the *Run DBCC CHECKDB* check box if you wish to verify the logical and physical integrity of all the objects in the specified database(s) after the restore completes.
- Click the **Restore** button.
- A confirmation box will appear. Click **OK** to restore or **Cancel** to cancel.

Database Ownership After Restore

This section does not apply to restores to the Original Location, as database ownership will not change.

The *amandabackup / CarboniteUser* user is used to initiate the restore operations for Microsoft SQL Server. It will thus become the new owner of the restored databases during restores to a different computer or different instance.

This is necessary for several reasons:

- It cannot be assumed that the user who previously owned the database exists during the restore process.
- The user performing the restore must have the necessary permissions to apply transaction log backups to the newly restored database.
- Likewise, the user performing the restore must have full and total access to the database, including ownership changes.

Since the *amandabackup / CarboniteUser* user is required be added as a Sysadmin, it will always satisfy these requirements.

After the restore, any system administrator can change the database ownership to any user. Please see this Microsoft Article for more information about changing database ownership: Changing the Database Owner

Restoring a Microsoft Exchange Server Database

Exchange databases can be restored to the same server or to an alternate server. Three types of recovery operations are allowed by Microsoft Exchange:

- Roll-forward
- Point in time
- Full restore

The state of the Microsoft Exchange server when the restore is started determines what type of recovery will be available. Please see this Microsoft documentation for details on Exchange recoveries.

All Exchange versions support restores to the Original location. In addition, you can restore Exchange 2007 backups to a Recovery Storage Group and Exchange 2010 or 2013 backups to a Recovery Database.

If you wish to restore an individual mailbox or message, please refer to Recover a deleted mailbox or a mailbox item below.

Additional Requirements

Microsoft Exchange Server restores have many additional requirements. Some requirements may vary between different versions of Exchange.

Requirements for All Exchange Restores

- The Microsoft Exchange Replication service must be started for Exchange 2010 and higher.

- The Exchange Server must be installed and running on the system.
- The following prerequisites must be installed for all versions of Exchange (except for Exchange 2003):
    - .NET 4.0 or higher
    - Powershell 2.0 or higher

Requirements for Exchange Restores to an Alternate Machine

- The Exchange version and service packs on the alternate machine must be the same as the original machine.
- Exchange must be installed with the same Organization and Administrative Group name as the original server.
- The storage groups and databases must already exist on the alternate server. They must also have the same names as the original storage groups or databases.
- Because you are restoring to an alternate "recovery server" that has a different set of log files, the signatures on the log files must match those of the original server.
    - To ensure that the log file signatures match, either rename the *E0x.log* file located in the Transaction log directory or enable the **Do not mount the database** option while creating a Mailbox or Public folder store.
- The Volume Shadow Copy Service must be enabled.
    - In Windows 2003 Small Business Server edition, the Exchange Writer is disabled by default. Please follow the instructions in this Microsoft knowledgebase article to enable the Exchange Writer.
- The Exchange VSS Writer must be stable.
    - Run the command **vssadmin list writers** in the Windows command prompt to check that the state of the Exchange Writer.
    - If it is not in a *Stable* state, restart the *Microsoft Exchange Information Store* Service.

Additional Requirements for Selective Restores in Exchange 2003 and 2007

It is possible to selectively restore individual Public Folder Databases and Mailbox Databases in Exchange 2003 and 2007, but there are further requirements that must be met.

Individual databases can be restored without these additional requirements on Exchange 2010 and higher.

- Selective restore of Public Folder Databases and Mailbox Databases can only be performed to the Original location for Exchange 2003.
- Selective restore to the Original Location or a Recovery Storage Group is possible for Exchange 2007.

- The Database Files and the associated Logs must be selected on the restore page. They will be listed separately. You must select both.
- The *Delete existing transaction logs before running the restore* method (discussed below) is not allowed for selective restores to the original location in Exchange 2003 or 2007.
    - A *No Loss Restore - Do not delete existing transaction logs* restore will be performed for all selective restores, even if **Delete existing transaction logs** is selected.

Special Options

There are several special options to consider when restoring a Microsoft Exchange Server.

Restore To

Microsoft Exchange Server has a unique set of restore options:

- Original Location
- Alternate Location
- Recovery Storage Group or Recovery Database

Each Restore Location may have requirements in addition to those listed above. Each Restore Location option is designed to fill a specific purpose.

Restore To: Original Location

If chosen, this will restore the selected database(s) to their original location. Restore to the Original location is most frequently used to fix a corrupted Exchange database or resolve issues with missing or damaged transaction logs.

Additional Requirements

- Any database that is to be restored must be in the **Dismounted** state.
    - All Mailbox Databases and Public Folder Databases within the target Storage Group must meet this requirement, even if you are only restoring a single database.
        - If one or more databases selected for restore are mounted, the restore will fail with the error *vss initialization failed*.
- The *This database can be overwritten by restore* option must be enabled via the Exchange System Manager (Exchange 2003) or Exchange Management Console (Exchange 2007 and higher).

Restore Method

You will be presented with two options when you restore to the Original Location:

No Loss Restore - Do Not Delete Existing Transaction Logs

- The existing transaction logs will not be deleted during the restore. Once the restore is complete, Exchange will perform a roll-forward recovery if all sequential transaction logs are found.
- Because Public Folder Databases participate in replication by design, the *No Loss Restore - Do not delete existing transaction logs* method is not supported for any restores that include Public Folder Databases; the **Delete existing transaction logs before running the restore** method must be chosen for any such database.

Delete Existing Transaction Logs Before Running the Restore

- This is a "Point in time" recovery method.
- All transaction logs that are currently present in the *Log directory* of the target Storage Group(s) will be deleted at the beginning of the restore process. This includes transaction logs for Mailbox Database(s), Public Folder Database(s), and/or Storage Group(s).

Procedure to Restore Exchange Server Data to the Original Location

- Navigate to the *Restore* page and choose a Restore Point.
- By default, all databases will be selected. You may individually select a Mailbox Store Database or Public Folder Database & its associated Logs from Restore Select.
- Click the **Edit** button next to *Review the restore settings*.
- Select **Original Location** from the *Restore To* dropdown menu.
- Select **Overwrite Original** from the *name conflict settings* dropdown menu.
- Choose the Restore Method you wish to use.
- Review your selections, then click **Restore** to start the restore process.

Restore To: Recovery Storage Group or Recovery Database

Microsoft Exchange Server allows restore of Exchange Mailbox Stores to a Recovery Storage Group in Exchange 2007 and to a Recovery Database in Exchange 2010 & 2013.

Recovery Storage Groups/Databases allow users to restore the Mailbox Store to a production Exchange Server without dismounting the existing Mailbox stores.

Only Exchange 2007 and up will support Recovery Storage Groups/Databases.

It is not possible to restore a Public Folder Database to a Recovery Storage Group/Database. This functionality is not supported by Exchange. Any attempt to do so will result in a failed restore with the error "Restore could not proceed as VSS initialization failed."

Further details about the failure can be found in the Windows Application Event Logs.

Additional Requirements

All Exchange-specific requirements must be met. There are additional requirements for restores to a Recovery Storage Group/Database:

- The Recovery Storage Group or Recovery Database must be created before the restore begins.
    - Use the Exchange Management Tool or the Exchange Management Shell to create the Recovery Storage Group/Database
- All the recovery databases in the Recovery Storage Group must be *dismounted*.
- All the recovery databases must have the "*This database can be overwritten by restore*" option enabled.
- The Recovery Storage Group or Recovery Database folder must be empty.

If you are restoring to a Recovery Storage Group/Database on an alternate machine:

- The name of the alternate machine must exactly match the name of the original machine.
- The recovery server must have the same Organization name, Storage Group name, and Mailbox Database name as the original server.

Procedure to Restore Exchange Server Data to an Alternate Exchange Server

- Navigate to the *Restore* page and choose a Restore Point.
- By default, all databases will be selected. You may individually select a Mailbox Store Database or Storage Group and its associated logs from *Restore Select*.
- Click the **Edit** button to review the restore settings.
- Select **Recovery Database** from the *Restore To* dropdown menu.
    - For Exchange 2007: enter the name of the Recovery Storage Group
    - For Exchange 2010 and 2013: enter the name of the Recovery Database
- Review your selections, then click **Restore** to start the restore process.

Recover a Deleted Mailbox or a Mailbox Item

CSSB now supports mailbox level backup and restore. For additional information about mailbox level backup, please refer to the following sections in the User Guide:

- Exchange Local Mailbox Backup
- Exchange Online hosted for Office 365 Backup

To restore mailbox items from an Exchange Local Mailbox backup or an Exchange Online Mailbox backup, please refer to the following Knowledge Base articles:

- Restoring Microsoft Exchange Local Mailbox
- Restoring Microsoft Exchange Online hosted for Office 365

For specific instructions on restoring individual mailboxes or items from an Exchange Database backup, please refer to one of our Knowledge Base articles below:

- Restoring mailbox items in Exchange 2007
- Restoring mailbox items in Exchange 2010
- Restoring mailbox items in Exchange 2013

Restoring a Microsoft SharePoint Server

CSSB can be used for content recovery, web application recovery, and disaster recovery of Microsoft SharePoint. Restores can be performed to the original system or an alternate system.

SharePoint includes a Microsoft SQL Server component.

- SharePoint includes a Microsoft SQL Server component. The MSSQL component must be running, unless the *rebuild system databases* option is checked. Please see below for more details.

Additional Requirements

There are additional requirements for all SharePoint restores.

Requirements for all SharePoint restore to the original machine

- All SharePoint databases must be in the *Normal* state prior to restore.
    - State can be verified from the SQL Management Studio. The *Properties* menu for each database will display its current state.
- The following services must be started:
    - Windows SharePoint VSS Writer
    - Windows SharePoint Services Tracing
    - The applicable *OSearch* and *SPSearch VSS Writer* services must be enabled. The name of the controlling service depends on the version of SharePoint installed. Examples below:
        - For SharePoint 2010:
            - The SPSearch VSS Writer is controlled by the *SharePoint Foundation Search V4* service.

        - The OSearch VSS Writer is controlled by the *SharePoint Server Search 14* service.
    - For SharePoint 2007:
        - The SPSearch VSS Writer is controlled by the *Windows SharePoint Services Search* service.
        - The OSearch VSS Writer is controlled by the *Office SharePoint Server Search* service.

Requirements for SharePoint restore to an alternate machine

- The new machine must be running the same version of SharePoint with the same Service Packs and the same embedded SQL Server that the original machine was running.
- The hostname of the new machine must be the same as the original machine.
- The instance name of the new server must exactly match the instance name from the original server.
- SharePoint must be installed to the same location on the disk as when the backup was run.
- All new databases and log file locations must also match those from the original configuration.

Special Options

There are several special options to consider when restoring a Microsoft SharePoint Server.

Rebuild System Databases

This option is intended to be used following a bare metal restore.

All SharePoint installations include a Microsoft SQL Server component. The MSSQL component must be running to perform a SharePoint restore.

However, the Microsoft SQL Server component cannot be started after a bare metal restore because the system databases are not included in the bare metal recovery process.

Place a check mark next to *Rebuild System Databases* to rebuild the system databases during restore. This option will replace the existing system databases with the system databases from the backup. Once complete, CSSB will attempt to start the MSSQL server.

Once MSSQL starts, CSSB will automatically restore the remaining SharePoint information. There is no need to run multiple restores.

Please see this Knowledge Base article for more information.

Restore To

SharePoint databases can be restored to their original location or to an alternate location.

Please note that CSSB does not support Roll-forward restores for SharePoint. If a restore is targeted to the original location, any changes made to the database after the backup was taken will be lost.

Restore To: Original location

The following occurs during a restore to the Original location:

- The following SharePoint services are stopped before the restore begins:
    - Windows SharePoint Services Administration
    - Windows SharePoint Services Search
    - Windows SharePoint Services Timer
    - Office SharePoint Server Search
    - IS Admin Service
- The selected SharePoint database(s) and log files are restored to their original location.
    - Index search files will be restored, if all databases and logs within the backup are selected for restore.
    - Individual content databases can be selected for restore but index search files will not be restored.
- CSSB calls the SharePoint VSS writer which automatically detaches and then reattaches each database to the farm.
- The services which were stopped before the restore operation are restarted.

Restore To: Alternate Location

No services are stopped or started before a restore to an alternate location. Nor are the databases automatically attached.

Instead, the databases and log files are simply restored to disk in the given location. Please import them manually into SharePoint.

Restoring Hyper-V

CSSB supports the following types of restores:

- All virtual machines may be restored using *Restore All*.
- Individual virtual machines may be restored using *Restore Select*.
- The Initial Store may be restored using *Restore Select* on Windows Server 2008 R2 and Windows Server 2012 R2.

    - The Initial Store cannot be restored in Windows Server 2012.
- Files and folders within the backed up virtual machines can be restored.
    - Use the Files/Folders restore option to restore the virtual hard disks. These can be mounted to the file system and browsed for individual files and folders.
    - CSSB will offer to mount the drives automatically.

Restored virtual machines will be automatically imported into Hyper-V in most configurations. When restoring one or more virtual machines, all associated virtual hard disks will be restored.

Additional Requirements

Additional Requirements for Restoring Virtual Machines

CSSB cannot overwrite an active, running virtual machine when restoring to the Original Location. If you attempt to do so, the restore will be interrupted with the following message:

The guest VM that you are trying to restore appears to be in a running state. It will be turned off during the restore.

Are you sure you want to continue?

Choose **OK** to turn off the VM and continue the restore. Select **Cancel** to abort the restore process.

The above message will not appear if the name of the Virtual Machine has changed since the time the backup was taken. CSSB will proceed directly with the restore and overwrite the existing VM.

Additional Requirements for Restoring Files and Folders

CSSB will offer to automatically mount a restored virtual hard disk file once the restore is complete. It will also ask if you want the mounted VHD to be displayed in Windows Explorer after the restore.

- The operating system of the restore machine must be equal to or higher than the operating system of the machine where the backup was taken.
    - For example, a virtual hard disk that originated on a Windows Server 2012 system cannot be mounted to a Windows Server 2008 R2 system.
- If every possible drive letter is already in use, the VHD cannot be automatically mounted. There must be free, non-reserved drive letters available.
- The virtual hard disk must be using one of the following file systems to be automatically mounted.
    - FAT
    - FAT32

    - NTFS
    - ReFS, with some restrictions
        - ReFS disks can only be mounted on Windows Server 2012 and higher with the ReFS role enabled.
        - Additionally, the operating system version must match between the original system and the system performing the restore.
- Dynamic disks cannot be mounted.

Managing VHD Files After the Restore

You can manage your VHD files from within CSSB after the restore is complete. Simply right-click the Restore operation in the Report page. You will be presented with three options:

- Mount VHD: Select this option to mount a VHD file. It is subject to the same requirements shown above.
- Dismount VHD: Select this option to dismount a VHD file that is currently mounted.
- Delete VHD: Select this option to delete the VHD file from disk.

CSSB will not dismount the vhd files automatically. You must dismount via the Report page, or do so manually using Disk Management or Powershell.

Further information about mounting and dismounting VHD files can be found in this Knowledge Base article.

Restoring a MySQL Server

MySQL databases can be restored to the original location or as a .sql file on disk that can be manually imported into any compatible MySQL Server.

Additional Requirements

MySQL restores have additional requirements.

- The MySQL user specified on the *Backup* page must have the following privileges to restore the MySQL databases:
    - CREATE, DROP, INDEX, SHUTDOWN, INSERT, ALTER, UPDATE, TRIGGER, SUPER, REPLICATION CLIENT, CREATE VIEW
- The MySQL client utilities (*mysqldump* and *mysql*) must be installed, and the MySQL client version must be compatible with MySQL Server.
- The MySQL Server must be running at the time of restoration.

Restore Views

There are two ways to view databases for restore: **SQL Files** and **Databases**, each with its own benefits.

SQL Files

The *SQL Files* view is the default option for MySQL restores. This option is used when Restore All is selected and is the default view shown when Restore Select is chosen.

An SQL File Restore will always include all databases that were backed up. Selective restore is not possible with the SQL Files restore view.

An SQL File restore will create a .sql file on disk. This .sql file must be manually imported into any MySQL installation using MySQL tools or the command line.

Databases

The *Databases* view allows for individual databases to be selected for restore. A *Databases* restore will always restore the databases to their Original location and will import the databases into MySQL automatically.

The Databases view is only visible when *Restore Select* is chosen.

- Navigate to the *Restore* page and select a restore point.
- Click **Restore Select**.
- Change to the *Databases* view in the upper right.
- Select the databases to be restored.
- Click **Restore**.

Restoring an Oracle Server

Oracle databases can be restored to the original location or to an alternate location. Oracle restores can be highly technical and require access to the SQL*Plus utility that is included with the Oracle database.

Additional Requirements

Oracle restores have additional requirements:

- SQL*Plus must be installed and accessible to the user performing the restore.
- The Oracle VSS Writer service must be started and functional.

Recovery of ARCHIVELOG Databases to the Original Location

To Restore All ARCHIVELOG Databases to the Original Location

- Ensure that the instance is not started.
- Navigate to the *Restore* page. Select a Restore Point.
- Select the database files to be restored. This includes the datafiles, server parameter file, etc.
    - If the redo logs are missing, you must also select the appropriate archived redo log files.

- Select **Restore to Original Location** from the dropdown menu.
    - CSSB will restore the data to the same location the data was backed up from.
- Choose **Overwrite Original** as your *name conflict settings*.
- Click **Restore** to start the restore process.
- Once the database is successfully restored, open SQL*Plus and run the following commands:

```
cmd> sqlplus /nolog
sql> connect sys as sysdba
sql> shutdown immediate
sql> startup mount
sql> recover database using backup controlfile until cancel;
```

Note: In cancel-based recovery, recovery continues by prompting you with the suggested filenames of archived redo log files. Recovery stops when you input CANCEL instead of a filename, or when all redo logs have been applied to the datafiles.

- Continue applying redo log files until the last log has been applied to the restored datafiles, then cancel recovery by executing the following command:

```
sql> CANCEL;
```

- To finish, enter the following SQL command string:

```
sql> alter database open resetlogs;
```

- To verify that the database is in the open state (read/write mode), use the following command:

```
sql> select name, open_mode from v$database;
```

To Restore a Single Oracle Datafile to the Original Location

Ensure that the database is either mounted or open. If the database is open, run SQL*Plus from the Command Prompt and connect to the database:

```
cmd> sqlplus sys as sysdba
```

To recover the datafile, the tablespace needs to be taken offline. In the following example, a tablespace labeled "TEST01" needs to be restored.

```
sql> ALTER TABLESPACE TEST01 OFFLINE IMMEDIATE;
```

- Navigate to the *Restore* page in CSSB. Select the backup run that you want to restore from.
- Select the datafile to be restored.

    - Our example will use a datafile named *TEST01.dbf*.
    - If the redo logs are missing, select the appropriate archived redo log files as well.
- Select **Restore to Original Location** from the dropdown menu.
    - CSSB will restore the data to the same location from where the backup was made.
- Choose **Overwrite Original** as your *name conflict settings*.
- Click **Restore**.
- Once the database is successfully restored, open the SQL*Plus and run the following commands:

```
cmd> sqlplus /nolog
sql> connect sys as sysdba
```

- If the restored data file is older than the redo log files, then a media recovery needs to be initiated using the following command:

```
sql> RECOVER TABLESPACE TEST01;
```

- You will see the following message:

```
Specify log: {=suggested | filename | AUTO | CANCEL}
```

- Input the appropriate log file location. Type **AUTO** if you want Oracle to automatically find the appropriate logs and apply it.

```
Specify log: {=suggested | filename | AUTO | CANCEL}
AUTO
Log applied.
Media recovery complete.
```

- Once you have recovered all the tablespaces that you have lost, bring the offline datafile or tablespace back online using the following command:

```
sql> ALTER TABLESPACE TEST01 ONLINE;
```

To Restore All ARCHIVELOG Control Files

Note: When Oracle database control files are backed up by Volume Shadow Copy Service, the backed up files are referred to as *Snapshot Control* files. For example, a control file by the name of "CONTROL01.CTL" is backed up under the name "SNCF[DATABASE_SID].ORA". In the following example, the [DATABASE_SID] is *ORCL*. In this case, the snapshot control file will have the name "*SNCFORCL.ORA*".

- Ensure that the database is in the *NOMOUNT* state or can be started in *NOMOUNT* state by the Oracle VSS writer. You do not have to shut down the database manually.

- Navigate to the *Restore* page on the CSSB user interface. Select the backup run that you want to restore from.
- Select the snapshot Control File. In the above example, you would select *SNCFORCL.ORA* (since the SID of the database is *ORCL*).
- Select **Restore to Original Location** from the dropdown menu.
    - CSSB will restore the data to the same location it was backed up from.
- Choose **Overwrite Original** as the *name conflict settings*.
- Click on **Restore**.

Note: CSSB automatically stops the Oracle instance while restoring and then starts it again when the restore operation is finished.

- Make three copies of the SNCF[DATABASE_SID].ORA file and label them **CONTROL01.CTL**, **CONTROL02.CTL**, and **CONTROL03.CTL**.
- Copy these three files to the installation directory of the control files.
    - The default location is *ORAHOME**oradata**[DATABASE_SID]*\. For the example above, copy the files to the *C:**app**Administrator**oradata**ORCL*\ directory.
- Run sqlplus on Command Prompt Window:

cmd> sqlplus sys as sysdba

- Shut down the database and startup:

sql> SHUTDOWN IMMEDIATE
sql> STARTUP MOUNT

- Run the following command to get the database in consistent state:

sql> RECOVER database using backup controlfile until cancel;

- Continue applying redo log files until the last log has been applied to the restored datafiles, then cancel recovery by executing the following command:

sql> CANCEL;

- Open the database with the *RESETLOGS* option:

sql> ALTER DATABASE OPEN RESETLOGS;

- The Oracle database is now recovered. To verify that the database is in the open state (read/write mode), use the following command:

sql> select name, open_mode from v$database;

Recovery of the Server Parameter File

- Navigate to the *Restore* page in the CSSB. Select the backup run that you want to restore.
- Select the Server parameter file (*spfile*) to be restored.
    - Usually, it is named *SPFILE[ORACLE_SID].ORA*.
        - For example, if the SID of the database is "ORCL", then the name of spfile will be "*SPFILEORCL.ORA*".
- Select **Restore to Original Location** from the dropdown menu.
    - CSSB will restore the data to the same location from where the backup was made.
- Choose **Overwrite Original** as your *name conflict settings*.
- Click on **Restore**.
- Once the restore operation is done, run SQL*Plus.

cmd> sqlplus sys as sysdba

- Shutdown the database and startup again.

sql> SHUTDOWN immediate
sql> STARTUP

- Check the database whether it's up and running

sql> SELECT name, open_mode from v$database;

Recovery of NOARCHIVELOG Databases

To Restore All NOARCHIVELOG Datafiles to the Original Location

- Navigate to the *Restore* page in the CSSB user interface. Select the backup run that you want to restore from.
- Select all the data files and control files. CSSB backs up one of the three control files (i.e. CONTROL01.CTL).
- Select **Restore to Original Location** from the dropdown menu.
    - CSSB will restore the data to the same location from where the backup was made.
- Choose **Overwrite Original** as your *name conflict settings*.

- Click **Restore**.
- Make two extra copies of the CONTROL01.CTL file and label them *CONTROL02.CTL* and *CONTROL03.CTL*. For your information, CONTROL01.CTL file is in the *ORAHOME\ORADATA\[DATABASE_SID]* directory.
- Run *sqlplus* in a command prompt to start the recovery:

```
cmd> sqlplus /nolog
sql> connect sys as sysdba
sql> shutdown immediate
sql> startup nomount
sql> ALTER DATABASE mount;
sql> RECOVER DATABASE;
sql> ALTER DATABASE open;
```

Restoring All NOARCHIVELOG Datafiles to an Alternate Location

When *Restore to Alternate Location* is chosen, CSSB will simply restore the selected Oracle files to disk. They must be imported manually into the Oracle database.

Restoring Exchange Local Mailboxes

Any number of mailboxes can be selected for restore, including individual mailboxes.

Mailbox contents can be restored to the original server, an alternate server, an Exchange Online server, or to disk.

For all options except restore to disk, mailbox contents can be restored to a mailbox with the same name or a different name as the original.

Restores to disk can be done as .pst files (one for each mailbox) or as files (.eml, .png, etc) that correspond to each item within the mailboxes restored. For more information about restoring as files, please see How to Use Restored Files after an Exchange Mailbox Restore to Disk.

Additional Requirements

- An administrator account must be specified within CSSB. The username must be in the User Principal Name (UPN) or email format, such as UserName@Domain. For example, Admin@company.com and username@company.onmicrosoft.com are both in the correct format.
    - The Exchange Impersonation right must be granted to this administrator account.
        - This right can be granted to an account from the Admin page in the Office 365 portal.
- The Exchange Discovery Management role must be granted to this administrator account.
    - This role can be granted to an account from the Admin page in the Office 365 portal.

- When restoring to a local Exchange server, a server URL must be specified. The system performing the restore must be able to reach and resolve the server URL.
- When restoring to an Exchange Online account, the system performing the restore must be able to connect to the Office 365 account. Connectivity can be verified from any browser.
- Restore to an Alternate Mailbox is allowed with three exceptions:
    - An alternate mailbox cannot be specified for restores to disk. The files will always restore to a folder with the same name as the original mailbox.
    - An alternate mailbox is not allowed when "Restore All" is selected, even if the backup contains just one mailbox.
    - An alternate mailbox is not allowed when more than one mailbox is chosen in "Restore Select".
- When restoring to .pst files, Microsoft Outlook must be installed on the system.

Restoring Exchange Online hosted for Office 365

Any number of mailboxes can be selected for restore, including individual mailboxes.

Mailbox contents can be restored to the original server, an alternate server, to a local Exchange server, or to disk.

For all options except restore to disk, mailbox contents can be restored to a mailbox with the same name as the original or to an alternate mailbox with a different name.

Restores to disk can be done as .pst files (one for each mailbox) or as files (.eml, .png, etc) that correspond to each item within the mailboxes restored. For more information about restoring as files, please see How to Use Restored Files after an Exchange Mailbox Restore to Disk.

Additional Requirements

- An administrator account must be specified within CSSB. This user must have the Global Administrator role. The username must be in the User Principal Name (UPN) or email format, such as UserName@Domain. For example, Admin@company.com and username@company.onmicrosoft.com are both in the correct format.
    - The *Exchange Impersonation* right must be granted to this administrator account.
        - This right can be granted to an account from the Admin page in the Office 365 portal.
- The *Exchange Discovery Management* role must be granted to this administrator account.
    - This role can be granted to an account from the Admin page in the Office 365 portal.
- When restoring to a local Exchange server, a server URL must be specified. The system performing the restore must be able to reach and resolve the server URL.

- When restoring to an Exchange Online account, the system performing the restore must be able to connect to the Office 365 account. Connectivity can be verified from any browser.
- Restore to an Alternate Mailbox is allowed with three exceptions:
    - An alternate mailbox cannot be specified for restores to disk. The files will always restore to a folder with the same name as the original mailbox.
    - An alternate mailbox is not allowed when "Restore All" is selected, even if the backup contains just one mailbox.
    - An alternate mailbox is not allowed when more than one mailbox is chosen in "Restore Select".
- When restoring to .pst files, Microsoft Outlook must be installed on the system.

Restoring a MailStore Archive

MailStore archives can be restored to any system, even one where MailStore is not installed. However, a valid MailStore installation is required to access and manage the restored MailStore archives.

Functionality

CSSB will restore the MailStore archives to the chosen location on disk.

MailStore archives cannot be restored to their Original Location. An alternate location must be specified.

Once restore is complete, the MailStore archives can be managed via the MailStore application.

Additional Requirements

- MailStore archives cannot be restored to a folder where another MailStore archive exists. Please restore MailStore archives to an empty folder.
- MailStore archives can be restored to any system, even one where MailStore is not installed.

Restoring a Bare Metal Image

A Bare Metal Image backup can be restored using two methods:

- *Bare Metal Restore*: A Bare Metal Image backup can be restored to any computer by booting from the recovery media. This includes the ability to restore a physical machine to a virtual machine or a virtual machine to a physical machine. This is a total recovery option. Any data on the existing system, if such data exists, will be removed.
- *File/Folder Level Restore*: Files and folders can be restored from the Bare Metal Image backups. This is not a total recovery option.

Please remember that Exchange, Microsoft SQL Server, and SharePoint databases are backed up in separate Bare Metal Suite backup sets. If you perform a full Bare Metal Restore, you must restore these backup sets separately after the Bare Metal Restore is complete.

Note: CSSB cannot send an email notification upon success or failure of any Bare Metal Image restore. This includes both complete Bare Metal recovery and file/folder restores.

Windows Activation Following a Complete Bare Metal Restore

Windows licenses are often tied to a specific hardware GUID. Performing a complete Bare Metal Restore to different hardware may cause Windows throw errors that it is not activated or is not genuine. This behavior is expected, because the new system has a different hardware GUID than the original.

Self-help options exist for Windows Activation. Please see this Microsoft Knowledge Base article. You may also call Microsoft Support for assistance. Please visit this Microsoft page to find the phone number for your country.

Carbonite Support cannot assist with Windows activation.

Additional Requirements

- All restores require that the backup images be available in a disk location. External drives and network locations are ideal.
    - Bare Metal Image backups cannot be restored directly from the cloud. If the backup images are not available on disk at all, they must be downloaded manually before the restore begins.
    - If the backup images are available on disk, but that location is not accessible from the restore machine, you must copy the backup images on disk to a new location before the restore begins.
- There must be at least 2GB of memory on the system.
    - This requirement increases with both the number of files in the backup and with the number of incremental backups to be restored.
    - Restore will fail with *out of memory* errors if there is insufficient memory. Restore cannot proceed until more memory is added to the system.
- You may only perform one restore at a time.
- During a complete Bare Metal restore, the size of the hard drive(s) must be equal to or larger than the drives on the original system.
- During a complete Bare Metal restore, each disk involved in the restore must have the same sector size as the corresponding disk on the original system.
    - For example: If the original system had a disk with 512 bytes per sector, the new system must also have a disk with 512 bytes per sector.
    - Incompatible disks will not be shown (and cannot be selected) for restore in the recovery wizard.

For additional information about performing a Bare Metal Image restore, and steps to follow after a complete restore, please refer to this Knowledge Base article.

Disaster Recovery

Disaster Recovery refers to restoring the system and data of a partially or completely failed computer.

Planning for Disaster Recovery

Disaster recovery can be performed via a Bare Metal Image backup or by a combination of File System, System State, and various application backups.

Bare Metal Image (BMI) backups provide the fastest, easiest, and most comprehensive form of disaster recovery. For additional information, please refer to the Knowledge Base articles for Creating a BMI backup and Restoring a BMI backup.

Recovery via File System, System State, and application backups is significantly more involved. This form of disaster recovery requires that all relevant data be backed up. This typically requires several different backup sets of different backup types.

If you plan to recover using File System, System State, and application backups it is recommended that you simply back up everything:

- Create a *File System* backup set that contains all drives on the system.
- Create a *System State* backup set to get the system's registry, boot files, Active Directory, and more.
- Create one backup set for each *database* or *application present on the system*, such as Microsoft Exchange or Hyper-V.
- Ensure that a backup is taken of all data stored on the network, such as file shares or NAS devices.
    - CSSB can back up network storage using the *File System* backup type. Please refer to our Knowledge Base article on How to Use a NAS Device with CSSB for more information.

If everything cannot be backed up for some reason, the following is the minimum required for disaster recovery:

- A *File System* backup set that includes:
    - The entire system drive.
        - On most Windows system, C: is the system drive.
        - Temporary Files can be excluded.
        - Windows/system files can be excluded.

        - Database or application installation directories, if not on the system drive.
            - Any database or application that needs to be restored must have all its files on disk backed up in addition to the databases themselves.
            - Database files, such as .mdf files for Microsoft SQL Server, that are backed up separately as part of a *Database* or *Application* backup can be excluded.
        - User data that is not on the system drive.
- A *System State* backup set.
- A backup set for each type of *Database* or *Application*, such as Microsoft Exchange or Hyper-V.

Performing Disaster Recovery

Bare Metal Image backups provide the fastest, easiest, and most comprehensive form of disaster recovery. For additional information, please refer to the Knowledge Base articles for Creating a BMI backup and Restoring a BMI backup.

For Disaster Recovery that involves File System, System State, and other application backup sets, please review the process below.

It is recommended to use identical hardware to ensure that the System State backups restore correctly. The greater the difference in hardware, the more likely that the System State restore will encounter problems. These problems can range from failed restores to a complete inability to boot the system after the restore is complete. Systems that are fundamentally different, especially with regards to storage devices, are significantly more likely to encounter a problem.

The following is a high-level overview of the disaster recovery process using File System, System State, and application backup sets. The details may vary based on the type of data you are trying to restore.

- Install the version of Windows that was on the original system, including Service Packs.
    - Ensure that the new system has the exact same hard drive partitions, Windows installation directory, and Host Name as the original server.
    - Do not place the system in a Windows Domain or Workgroup.
- Log into the machine as an Administrator.
- Install CSSB and import your account's cloud certificate.
- Import existing backup sets and restart CSSB when it is complete.
- Choose the *File System* backup set that contains your system drive.
- Go to the *Restore* page.

- Choose the restore point.
- Choose to restore to the **Original Location** and **Overwrite Existing** files.
- Begin the restore.
    - Restore of some system files may fail because the system is in use. This is normal and should be expected.
- Restore the File System backup set that houses your application installation directories (if they exist and were separate from the System Drive backup set).
    - Choose to restore to the **Original Location** with **Overwrite Existing**.
- Restore the *System State* backup to the **Original Location** with **Overwrite Existing**.
- Reboot the system after the System State restore is complete.
    - Do not reboot until after the System State restore is complete. The Windows installation may become corrupted if the system is rebooted after the File System restore is complete but before the System State restore is complete.
        - If a reboot is performed at the wrong time, it is likely that the system will need to be formatted and the Disaster Recovery process restarted from the beginning.
- Once the server reboots, check to ensure that all database servers and applications are installed and in the running state (if they exist). Repair if necessary.
- Restore the *Database* or *Application* backups (if they exist).

An example walkthrough of a complete system recovery, including File System and System State, can be found in our Knowledge Base article on Recovery of a Windows 7 Computer.

Monitor Backups

The *Monitor* page displays the status of backups, restores, uploads and downloads. The most recent operation is displayed.

Visit the *Report* page for a complete history and fine control over individual backup runs.

All operations can be cancelled. Some can be paused and resumed, as shown on the following chart:

Operation	Cancel	Pause	Resume
Backup to Cloud	Yes	Yes	Yes
Backup to Disk	Yes	No	No

Upload to Cloud	Yes	Yes	Yes
Download	Yes	No	No
Restore	Yes	No	No

Canceling an upload will also cancel all queued uploads in the same backup set.

When the **Resume** button on the *Monitor* page is clicked, CSSB will resume all uploads for the current backup set that are currently paused or failed. To resume one specific job, locate it on the *Report* page.

Upload speeds are shown on the Monitor page. The current speed and the average speed are displayed in the details column.

Reports and Backup History

The *Report* page display contains the history of all operations. Backup runs and other operations can be managed from this page individually or in groups.

Report History

A history of all operations will be displayed on the Report Page. Each row references a specific backup run or other operation and contains many details.

Each backup run will have an icon indicating its status.

- A green checkmark icon indicates the operation was **successful**.
- A red X icon indicates the operation **failed**. Please see the Details column for more information.
- A yellow warning icon indicates a **warning**. The operation was successful with a minor issue or failed in an expected way. Check the note in the *Details* column for more information.
- A grey trash can indicates that the backup was **purged** or deleted. The details column will contain the date the backup was deleted.
- Two grey vertical bars indicates that an upload is **paused**. It can be resumed by selecting the *Resume* option from the right-click menu.
- A blue circle made of dots indicates that an operation is **in progress**.
- A clock indicates that an upload is **queued**. Queued uploads will be uploaded in the order they were created.

Blue info icons may appear in the *Details* column for any operation and can be clicked to obtain further information.

Manage Report History

You can click on any column heading to sort by that field. Columns can be dragged and dropped to display in any order.

By default, the report history is kept forever. All operations and backup runs will be displayed. Visit the **Preferences > Advanced** menu option to change this value.

Report Templates

CSSB allows you to customize your historical view. Click on the **Define Template** button, specify the template name, and choose the columns you want to include. Click on **Save Template** to add this new template.

To load a specific template, you can select it from the dropdown next to the **Define Template** button.

Click the **Export Report** button to export your loaded template view to a Comma Separated Values (CSV) file. This file can be printed or imported into other applications to create spreadsheets and charts.

Toolbar and Right-Click Menu

Right-click any backup run or operation to get a list of tasks that can be performed. These tasks are also available in the toolbar at the top of the *Report* page.

Delete

This option allows you to delete data associated with backup runs and other operations.

- **Delete Backup Data from Disk:** The data will be deleted from the disk.
- **Delete Backup Data from Cloud:** The data will be deleted from the cloud. Deletion of data from the cloud can take a long time.
- **Delete Backup Run:** The data will be deleted from both disk and cloud. Additionally, the backup run itself will be deleted from the Report page.

When you delete a cloud backup, its status in the Carbonite Portal will change to *Marked for Deletion* until the deletion is complete.

Upload

This option allows you to initiate an upload of a backup on disk or resume a paused upload.

Occasionally, failed uploads are unable to be restarted. The *Upload* options will be disabled if an upload is not possible. If the *Upload* options are disabled, please begin a new backup.

Change Retention Period

Select a backup and choose this option to change its retention period. This option only works for backups using time-based retention.

Verify Backup Data

This option checks that all backup files are present and available. It does not check the integrity of the data.

Select a Backup to Disk to verify a backup on disk. Select a Backup to Cloud or Upload from Disk to verify the data on the cloud.

The larger the backup, the longer verification will take. Verification of backups on the cloud can take a long time.

Verification cannot be performed for *in-progress*, *paused*, *purged*, or *queued* backups or uploads. It is also not present for Backup to Cloud operations that are *failed* or *cancelled*.

Information Column

The column on the right side of the *Report* page includes additional information about the selected backup run.

The **Show Contents of Backup Run** button loads a list of everything inside the selected backup run. The larger the backup, the longer it takes for the list to populate.

The list is paginated. Use the selector below the list to change to a different page.

Click the **Prepare to Restore** button to be taken to the **Restore** page. Only the chosen backup run will be selected for restore.

Administration

The *Cloud*, *Tools*, and *Preferences* menus are found at the top of the screen. Each contains additional options to help administrate backups.

Cloud Menu

The *Cloud* menu contains options related directly to the cloud storage associated with your account.

Import Cloud Certificate

Use this option to import a cloud certificate at any time.

Check Cloud Connection

This option performs a quick test to see if the system can connect to the cloud.

Tools Menu

The Tools menu option contains tools related to CSSB operation that do not fit anywhere else.

Restart Background Service

This tool restarts the two background services: Carbonite Safe Server Cloud Controller and Carbonite Safe Server Backup Controller. These services must be restarted after certain advanced configuration changes, particularly those done in the Preferences > Advanced menu.

Import Existing Backup Sets

CSSB stores special metadata for every backup run. This metadata is known as the *Backup Catalog*. The metadata keeps track of when a backup was taken, what is inside each backup, the order in which incremental backups should be restored, and more. A restore cannot proceed without the catalog information.

Using this tool is the first step when restoring to a new machine. Please refer to our Knowledge Base article on Importing Existing Backup Sets for more information.

Move Local Backups

This tool moves local backups from one location to another on disk. If local backups have been moved manually in Windows Explorer, you can also use this function to update the location in CSSB. Please refer to our Knowledge Base article on Moving Local Backups to a New Folder for more information.

Check Dependencies

Exchange backups have additional dependencies above and beyond those required for standard CSSB operation. This tool checks if those dependencies are installed. It also displays the version of the Java Runtime Environment installed on the system.

Network Location

By default, most CSSB operations run as the *amandabackup / CarboniteUser* account. This includes access to network storage.

In the *Network Location* menu, you can tell CSSB to use a different user account to access network storage, if you so choose.

Step 1: Choose if you want to use the *amandabackup / CarboniteUser* account or if you wish to connect using other user credentials.

Step 2: If you wish to use other credentials, fill in the UNC path of the location you plan to use for storage and enter your username / password.

Step 3: Place a checkmark in the *Verify destination access rights* option if you want CSSB to make sure the account you specified in Step 2 has the correct access rights.

Step 4: Click **Save**.

Preferences Menu

The *Preferences* menu allows users to set various global settings. Any value that is set in the Preferences menu will apply for all backup sets. If a value is set globally in the Preferences menu and set in a backup set, the backup set value will be used.

Bandwidth Management

This option allows users to limit the amount of bandwidth used by CSSB for all backup sets.

The value chosen here is not divided among backup sets. It will apply to each backup set separately. For example, each backup set will have a limit of 500kbps if a value of 500kbps is chosen in Bandwidth Management.

Email

Notifications can be sent on success or failure of any given operation as described in the **Sets** section of this User Guide.

A SMTP server must be specified to receive notifications. The following information must be provided by the user:

- *From:* Please specify the email address that will be in the From field of each email notification. Be sure to whitelist this address so the notifications are not flagged as spam or junk.
- *To:* Input the email addresses to which the notifications will be sent. Separate multiple email addresses with a "," (comma) character. In CSSB v4.14 and higher, you can separate multiple email addresses with a "," (comma) or ";" (semi-colon) character.
- *Outgoing SMTP Server:* This is the IP address or hostname of the outgoing SMTP server.
- *Port:* This is the SMTP server port.
- *Security:* Choose the type of security used for the email notifications.
    - The proper value must be chosen, and the SMTP server must support the chosen value.
    - For example, if you are using Exchange server, please enable **Require TLS encryption** option on the Exchange server. To configure Transport Layer Security Encryption for clients, please see the Microsoft KB article titled: How to help protect SMTP communication by using the Transport Layer Security protocol in Exchange Server
- *Username:* This is the email address/username which is to be used to log into the mail server.
- *Password:* This is the password for the above email address/username.
- *Authenticate:* Enable if the SMTP server requires authentication. Disable if it does not.

Use the **Send Test Mail** button to send a test mail. Review the SMTP server information provided or the settings on the SMTP server itself if the test fails.

The *Notification Preferences* option allows users to choose if notifications will be sent on Failure, on Success, or both.

Advanced

The *Advanced* menu contains a variety of optional, advanced features.

Data Transfer and Networking

- *Proxy Servers:* A proxy server can be used for data transfer to and from the cloud.
    - *Proxy Host:* Enter the proxy server's host address.
    - *Proxy Port:* Enter the port used for access to the proxy server.

Please refer to our Knowledge Base article on Configuring Carbonite to Work With Authenticated Proxy Servers for more information.

- *Ports:* Ports 10080 and 10081 are used internally for transfer of data during backup or restore. These ports can be changed if necessary. These ports have nothing to do with firewall configuration, and do not need to be opened in a firewall.
- *Backup Port:* Enter the port used to transfer data during a backup.
- *Restore Port:* Enter the port used to transfer data during a restore.

The background services must be restarted if the ports are changed. Click **Tools > Restart Background Service** to restart the services.

- *Threads:* Multiple connections are opened with the cloud servers during upload or download of data. Each such connection is called a *thread*. Increasing the number of threads can increase data transfer speed on some networks. However, using more threads than necessary can cause slower data transfer. The default number of threads is **3** (three) and the maximum number is **10**.
    - *Max. Download Threads:* Use this setting to change the number of threads used by CSSB when downloading data from the cloud.
    - *Max. Upload Threads:* Use this setting to change the number of threads used by CSSB when uploading data to the cloud.

Any change to the number of threads will not affect any data transfers currently in progress.

Retention Policy Enforcement

Retention policies are checked once per day. Those that have expired are purged.

- *Purge old backup data every day at:* Retention policies will be checked at the time specified here.

The *Report* page shows old backups, even those that have been purged.

- *Keep expired backup run reports for:* Set the length of time CSSB will display reports of old backup runs. Only backup runs that have expired or failed will be removed. Active backups will not be removed.

General

Several folders are used by the program. You can change the locations of these folders.

- *Temporary Directory:* Temporary files created during backups are stored in this directory. This folder must be on a local drive.
    - The background services must be restarted to change the Temporary Directory. Click **Tools > Restart Background Service** after changes are saved.
- *Backup Directory:* Backups are saved to this directory in a folder with the name of the backup set by default.
    - Changes made to the Backup Directory setting will only apply to new backup sets. It will not change the location for existing backup sets.
    - Please refer to *How to move local backups to a new folder* if local backups must be moved to a new location.
- *Download Directory:* During a restore, backup archives are saved to this directory by default.
- *Restore Directory:* When restoring to an alternate location, files are saved to this directory by default.

amandabackup / CarboniteUser User

A user named *amandabackup / CarboniteUser* is created during installation. CSSB stores the password chosen for *amandabackup / CarboniteUser*.

If this password changes in Windows, it must be updated manually in CSSB. Click the **Update Password** button to update the stored password for the *amandabackup / CarboniteUser* user.

Please refer to our Knowledge Base article on Password for more information.

Language

Carbonite Safe Server Backup (CSSB) has full support for two languages and limited support for eight more.

Languages with Full Support

The following languages are fully supported by Carbonite Safe Server Backup.

- English
- German

Languages with Limited Support

The following languages have limited support by Carbonite Safe Server Backup. Translations for languages with limited support may be incomplete.

- Japanese
- Simplified-Chinese

- Traditional-Chinese
- Korean
- Thai
- Italian
- Portuguese
- Spanish

The language that you choose in CSSB must match the language used for Windows on the system. For example, you should choose German as the language in CSSB if you are running a German version of Windows with the German language pack.

File names that contain characters cannot be backed up unless Windows includes those characters. For example, even if CSSB is set to the Japanese language, it cannot back up files that contain Japanese characters unless Windows is running in Japanese.

Note: Bare Metal Image backups are only supported in the English and German languages. Systems that run any other language may experience failures when attempting to perform a Bare Metal Image backup or restore.

Help and Support

Need help? Have questions? Don't worry! We have several self-help options and a support team.

Knowledge Base

There is a wealth of information at your disposal in the Carbonite Safe Server Backup Knowledge Base. All our documentation is contained within and is frequently updated.

If you hit a snag during setup, check the **Getting Started** section of this User Guide for some helpful links.

If you get an error in Carbonite Safe Server Backup, search for it in the Knowledge Base. Articles are added and updated frequently.

Support

Carbonite offers technical support through our *Carbonite Customer Care* team. Our team is trained on the CSSB product and ready to help you troubleshoot any issues.

Contact Us

You can contact Carbonite Customer Care for help with Carbonite Safe Server Backup by phone or email. Please refer to our Contact Us page for details.

You can also open your default web browser to our Contact Us page by clicking **Help > Contact Support** in the CSSB menu.

Subscription Details

Click **Help > Show Subscription Details** to see information about your account. This information includes which plan you have purchased, your expiration date, your current cloud storage quota, and your Carbonite Account ID.

Collect Log Files

Carbonite Customer Care may request logs from your system if more technical information is required to accurately troubleshoot your issue. In most cases, this can be done from the CSSB user interface. Please refer to our Knowledge Base article on Collecting Log Files for Troubleshooting for more information.