



Carbonite Server Backup Image Plug-in 8.3

User Guide



© 2017 Carbonite, Inc. All rights reserved.

Carbonite makes no representations or warranties with respect to the contents hereof and specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. Furthermore, Carbonite reserves the right to revise this publication and to make changes from time to time in the content hereof without obligation of Carbonite to notify any person of such revision of changes. All companies, names and data used in examples herein are fictitious unless otherwise noted.

No part of this document may be reproduced, transmitted, transcribed, stored in a retrieval System or translated into any language including computer language, in any form or by any means electronic, mechanic, magnetic, optical, chemical or otherwise without prior written permission of:

Carbonite, Inc.
Two Avenue de Lafayette
Boston, MA 02111
www.carbonite.com

Carbonite and the Carbonite logo are registered trademarks of Carbonite, Inc. Product names that include the Carbonite mark are trademarks of Carbonite, Inc. All other products or company names mentioned in this document are trademarks or registered trademarks of their respective owners.

Acknowledgements: Two encryption methods, DES and TripleDES, include cryptographic software written by Eric Young. The Windows versions of these algorithms also include software written by Tim Hudson. Bruce Schneier designed Blowfish encryption.

“Part of the software embedded in this product is gSOAP software. Portions created by gSOAP are Copyright 2001-2006 Robert A. van Engelen, Genivia Inc. All Rights Reserved. THE SOFTWARE IN THIS PRODUCT WAS IN PART PROVIDED BY GENIVIA INC AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.”

The Carbonite Server Backup Agent, Carbonite Server Backup CentralControl, and Carbonite Server Backup Director applications have the encryption option of AES (Advanced Encryption Standard). Advanced Encryption Standard algorithm (named Rijndael, pronounced “Rain Doll”) was developed by cryptographers Dr. Joan Daemen and Dr. Vincent Rijmen. This algorithm was chosen by the National Institute of Standards and Technology (NIST) of the U.S. Department of Commerce to be the new Federal Information Processing Standard (FIPS).

The Carbonite Server Backup Agents and Carbonite Server Backup Director applications also have the added security feature of an over the wire encryption method.

Document History

Version	Date	Description
1	October 2017	Carbonite Server Backup-branded Image Plug-in 8.3x guide.

Contents

1	Introduction to the Image Plug-in	4
2	Add an Image backup job.....	6
2.1	Protect SQL Server databases in AlwaysOn Availability Groups	8
2.2	Encryption settings.....	9
3	Run and schedule backups and synchronizations.....	11
3.1	Schedule a backup.....	12
3.2	Run an ad-hoc backup.....	15
3.3	Synchronize a job	16
4	Restore from Image Plug-in backups	17
4.1	Restore Windows volumes from an Image backup.....	17
4.2	Restore files and folders from an Image backup.....	19
5	Monitor computers and processes	24
5.1	View computer and job status information	24
5.2	View current process information for a job	25
5.3	View a job's process logs and safeset information	27
5.4	View and export recent backup statuses	29
6	Carbonite Server Backup Support	30
6.1	Contacting Carbonite.....	30

1 Introduction to the Image Plug-in

To back up Windows volumes as images, install the Image Plug-in with the 64-bit Windows Agent. Unlike the Windows Agent, which enumerates and backs up individual files and folders during a backup, the Image Plug-in sequentially backs up all blocks on a volume. Because backups with the Image Plug-in require significantly less processing than backups with the Windows Agent, the time required for a backup can be significantly reduced.

After the first “seed” backup of a volume, in which all data from the volume is sent to the vault, the Image Plug-in uses Changed Block Tracking to determine which blocks have changed. In subsequent Image backups, the Plug-in only reads and backs up changed blocks to the vault.

When creating an Image backup job, you can select specific volumes to back up, or create a Bare Metal Restore (BMR) job that backs up all volumes, partitions, and data required for restoring a system to new hardware. You can also back up data on Windows storage spaces. See [Add an Image backup job](#).

Note: The Image Plug-in does not back up or restore the configuration of Windows storage spaces. In a disaster recovery, you can configure storage spaces manually, and then restore volumes to the storage spaces.

You can restore entire volumes and specific files and folders from Image backups. See [Restore Windows volumes from an Image backup](#) and [Restore files and folders from an Image backup](#). You can also use the System Restore application to restore systems from Image Plug-in BMR backups to new hardware. For more information, see the System Restore *User Guide*.

Using Image Plug-in version 7.5 or later, you can back up volumes with SQL Server database files. This option creates application-consistent database backups, so that separate SQL Server Plug-in jobs are not required. You can then mount these safesets, and restore database files from the backups. See [Add an Image backup job](#) and [Restore SQL Server database files or objects from Image backups](#).

You can use the Image Plug-in only on supported 64-bit Windows operating systems with the NTFS file system. The Image Plug-in is not supported with ReFS file systems. The Plug-in supports both UEFI and BIOS, and MBR and GPT disks. For a complete list of supported platforms, see the Windows Agent release notes.

Image Plug-in Feature Summary

- Backs up volumes as images, which significantly reduces the amount of time required for a backup
- Backs up volumes from UEFI-based or BIOS-based systems. Restores volumes from UEFI-based system backups to UEFI-based systems, and restores volumes from BIOS-based system backups to UEFI-based or BIOS-based systems.
- Backs up system volumes, data volumes, or both in a single job
- Backs up and restores data on Windows storage spaces.

Note: The Image Plug-in does not back up or restore the configuration of Windows storage spaces. In a disaster recovery, you can configure storage spaces manually, and then restore volumes to the storage spaces.

- Restores entire volumes to live volumes, or mounts a safeset so you can restore specific files and folders
- Creates Bare Metal backups that can be restored using System Restore
- Managed using Portal

Note: You cannot manage Image Plug-in backups and restores using the legacy Windows CentralControl.

For more information, see the Portal help and the Windows Agent release notes.

2 Add an Image backup job

After a Windows computer with the Image Plug-in is added and configured in Portal, you can create an Image backup job. The Image Plug-in sequentially backs up all blocks on a volume instead of backing up specific files and folders. Because this process can require less processing than a traditional Windows backup job, the backup time can be significantly reduced.

In an Image backup job, you can select the following options:

- Specific volumes to back up
- Bare Metal Restore (BMR). This option backs up volumes that are needed to boot up the system after a system recovery. A BMR backup includes the volume where the operating system is installed, and the EFI system partition (ESP) on a UEFI-based system or the volume with the master boot record (MBR) on a BIOS-based system. In a disaster recovery situation, you can use the System Restore application to restore systems from BMR backups.

Note: BMR backup jobs can also be created using the Windows Agent without the Image Plug-in. Regardless of how a BMR backup was created, you can restore the backup using the System Restore application.

- Volumes with SQL Server database files. This option creates application-consistent SQL Server database backups, so that separate SQL Server Plug-in jobs are not required. Image Plug-in version 7.5 or later is required for this functionality.

After creating an Image backup job, you can run the job manually, and schedule the backup job to run. See [Run and schedule backups and synchronizations](#).

In a seed backup, the Image Plug-in processes data for every block on a volume— even blocks that are empty. The amount of data reported in the backup log for a seed backup could be larger than the amount of data actually on the volumes.

Note: If the Image Plug-in was installed silently, the machine must be restarted before Changed Block Tracking (CBT) can identify data that has changed since a previous backup. Without CBT, the Agent reads all data on a volume before backing up the changed blocks.

To add an Image backup job:

1. On the navigation bar, click **Computers**.

The Computers page shows registered computers.

2. Find a Windows computer with the Image Plug-in, and click the computer row to expand its view.

If a backup job has not been created for the computer, the system can attempt to create a backup job automatically. However, this job will not use the Image Plug-in. To create an Image backup job, click **Configure Manually**. To automatically create a default Windows Agent backup job for the computer, click **Auto Configure**.

3. Click the **Jobs** tab.

If a valid vault connection is not available for the computer, you cannot access the **Jobs** tab. You must add a vault connection before you can create a job. For more information, see the Portal online help.

4. In the **Select Job Task** menu, click **Create New Image Job**.
5. In the **Create New Job** dialog box, specify the following information:
 - In the **Name** box, type a name for the backup job.
 - In the **Description** box, optionally type a description for the backup job.
 - In the **Destination** list, select the vault where you want to save the backup data.
A vault only appears in the list if it assigned to the user, or if the user added it on the computer's Vault Settings tab.
 - In the **Log File Options** list, select the level of detail for job logging.
Note: For Image Plug-in jobs, the selected logging level does not affect the content of the logs.
 - For new backup jobs, the encryption method is AES 256 bit. Existing jobs can have other encryption methods. See [Encryption settings](#).
 - In the **Password** and **Confirm Password** boxes, enter an encryption password. You can also enter a password hint in the **Password Hint** box.

Note: For Image Plug-in jobs, the selected logging level does not affect the content of the logs.

6. In the **Select Volumes for Backup** box, do one of the following until the **Backup Set** box shows the volumes that you want to back up:
 - To back up specific volumes, select the check box for each volume that you want to back up, and then click **Include**.

- To back up volumes that are needed to boot up the system after a system recovery, select the **Bare Metal Restore** check box, and then click **Include**.

Note: In addition to restoring systems from Bare Metal Restore (BMR) backups using the System Restore application, you can restore specific volumes, files, and folders from BMR backup jobs created using the Image Plug-in.

- To back up volumes with SQL Server database files, and create application-consistent SQL Server database backups, click **Application Aware Backup**, and then select the **SQL Volumes Protected** check box.

7. Click **Create Job**.

The job is created, and the **View/Add Schedule** dialog box appears. You can now create a schedule for running the backup. Click **Cancel** if you do not want to create a schedule at this time.

For information about how to run and schedule the backup job, see [Run and schedule backups and synchronizations](#).

2.1 Protect SQL Server databases in AlwaysOn Availability Groups

You can protect SQL Server databases in AlwaysOn Availability Groups using the Windows Agent and SQL Server Plug-in, or the Windows Agent and Image Plug-in version 7.5 or later.

If you back up a database in a secondary replica, a copy-only backup of the database is performed. Copy-only backups do not affect the sequence of conventional SQL Server backups. Microsoft only supports copy-only backups of secondary databases (see <http://msdn.microsoft.com/en-us/library/hh245119.aspx>).

Note: If a backup job includes secondary databases and databases that are not in a secondary replica, a copy-only backup will be performed for all databases in the job. Do not include a secondary database in the same job as a standalone database.

To protect SQL Server databases in AlwaysOn Availability Groups, do one of the following:

- Install the Windows Agent and plug-in on the server where the primary replica is hosted.

If you use the SQL Server Plug-in, you can run a full backup of the primary databases, followed by full or transaction log backups. If the primary replica becomes a secondary replica after a failover, the Agent automatically runs copy-only database backups instead of full backups. Transaction log backups remain the same.

If you use Image Plug-in version 7.5 or later, you can run application-consistent image backups of the volumes with database files. If the primary replica becomes a secondary replica after a failover, the Agent automatically runs copy-only backups.

- Install the Windows Agent and plug-in on a server where a secondary replica is hosted. This backup strategy offloads backup processing to a non-primary server.

If you use the SQL Server Plug-in, you can run a copy-only backup of the secondary database, followed by copy-only or transaction log backups. If the secondary replica becomes the primary

replica after a failover, the Agent automatically runs full backups instead of copy-only backups. Transaction log backups remain the same.

If you use Image Plug-in version 7.5 or later, you can run application-consistent image backups of the volumes with secondary database files. The Agent automatically runs copy-only backups of secondary database files. If the secondary replica becomes the primary replica after a failover, the Agent automatically runs full backups instead of copy-only backups.

Note: If the availability mode of the secondary replica is asynchronous-commit, transaction logs on the secondary database could lag behind the primary replica database. If the secondary database is being backed up, data loss could occur.

- Install the Windows Agent and plug-in on the primary replica server and on secondary replica servers. This strategy ensures that backups continue even if one of the replicas is down.

If you use the SQL Server Plug-in, you can run a full backup on the primary replica, followed by full or transaction log backups. You can also run copy-only backups on the secondary replicas, followed by copy-only or transaction log backups.

If you use Image Plug-in version 7.5 or later, you can run application-consistent image backups on both the primary replica server and the secondary replica server. The Agent automatically runs copy-only backups of secondary databases.

If a SQL database in an AlwaysOn Availability Group is hosted on a SQL Server Failover Cluster Instance, install the Agent, SQL Server Plug-in and Cluster Plug-in on each physical node, and configure jobs on the virtual node. Full backups will run if the database is a primary database, and copy-only backups will run if the database is a secondary database.

Note: Only GPT disks are supported for Image backups (including application-consistent Image backups of volumes with database files) in a cluster.

For information about restoring SQL Server databases in AlwaysOn Availability Groups, see [Restore databases in AlwaysOn Availability Groups](#).

2.2 Encryption settings

Encryption settings specify the encryption type for backup data at rest on the vault. AES 256 bit encryption is the only encryption type available for new backup jobs.

If an existing job uses another encryption type (e.g., AES 128 bit, Blowfish, DES, Triple DES, None), you can continue to encrypt the job using that type. However, if you change the encryption type for an existing job, you cannot change the encryption type back to the original type. Only AES 256 bit will be available.

If you change encryption options for an existing job, it will force a new full backup (i.e., a reseed). The next backup will take longer than previous delta backups, and the amount of data stored on the vault will increase temporarily, depending on your retention settings.

Encryption password

You must enter a password for the encrypted backup data. The password is case-sensitive. To recover the data, you must provide the encryption password that was entered when the files were backed up.

You can also enter a password hint. When restoring data, you can view the password hint to remind you of the encryption password for this job.

Important: If you forget the encryption password, you lose access to the data. You cannot retrieve the password from the system.

3 Run and schedule backups and synchronizations

After a backup job is created, you can run it manually (ad-hoc) at any time and schedule it to run.

When running or scheduling a backup, you can specify the following settings:

- **Retention type.** The retention type specifies the number of days a backup is kept on the vault, how many copies of a backup are stored online, and how long backup data is stored offline.
- **Deferring.** You can use deferring to prevent large backups from running at peak network times. When deferring is enabled, the backup job does not back up any new data after the specified amount of time and commits the safeset to the vault, even if some data in the job is not backed up. Changes to data that was previously backed up will be backed up, regardless of the specified amount of time.

When the job runs again, the Agent checks for changes in data that was previously backed up, backs up those changes, and then backs up the remaining data.

If an Image Plug-in backup job is deferred while a volume is being backed up, the backup for the volume is incomplete and data from the volume cannot be restored. However, you can restore volumes, and files and folders from volumes that were completely backed up in the job before the job was deferred.

Note: Backups to SSI files on disk cannot be deferred.

- For an application-aware Image Plug-in job, you can specify whether to truncate SQL Server database transaction logs.

Note: If you also back up databases using another tool (e.g., native SQL Server backup), use only one tool for truncating logs.

When you schedule a job to run, you can also set the compression level for the data. The compression level optimizes the volume of data sent to the vault against the speed of processing. The default compression level is usually the optimal setting.

When a backup job first runs, all data selected in the job is backed up to the vault. This initial backup is called a “seed” backup. In subsequent backups, only data that has changed is backed up to the vault, unless a reseed is required (e.g., after a job’s encryption password has changed). In a reseed, all data selected in a backup job is sent to the vault again, even though it has already been backed up.

After running a backup, you can view logs to check whether the backup completed successfully.

In some cases, you must synchronize a backup job before you run it or restore data from the job. When you synchronize a job, the Agent checks which safesets for the job are online and available for restore. See [Synchronize a job](#).

3.1 Schedule a backup

After creating a backup job, you can add one or more schedules for running the job automatically.

You can create complex schedules for a job by creating multiple schedules. For example, you can schedule a backup job to run at midnight every Friday, and schedule the job to run at 8 pm on the first day of every month.

If a job is scheduled to start at exactly the same time by multiple schedules, the job only runs once at the scheduled time. If the jobs have different retention types, the retention type of the schedule that is highest in the list is applied to the resulting safeset. For example, in the following screenshot, the job is scheduled to run at 12 AM each Saturday with the Weekly retention type, and every day at 12 AM with the Daily retention type. On Saturdays, the job runs only once at 12 AM. Because the schedule with the Weekly retention type is higher in the list than the schedule with the Daily retention type, the Weekly retention type is applied to the safeset.

Note: If a job is scheduled to run at slightly different times, the Agent attempts to run the job according to each schedule. For example, if a job is scheduled to run at 11 PM by one schedule and 11:01 PM by another schedule, the Agent will attempt to run the job twice. Try to avoid overlapping schedules; problems can occur if a job is scheduled to run twice in a short period of time.

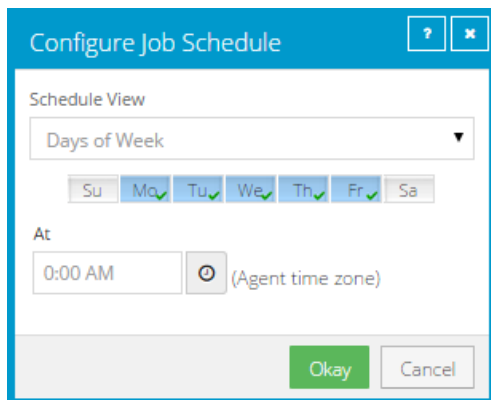
To schedule a backup:

1. Do one of the following:
 - On the navigation bar, click **Computers**. Find the computer with the backup job that you want to schedule, and click the computer row to expand its view. On the **Jobs** tab, find the job that you want to schedule. In its **Select Action** menu, click **View/Add Schedule**.
 - Create a new backup job. The **View/Add Schedule** dialog box appears when you save the job.
2. In the **View/Add Schedule** dialog box, click **Add Schedule**.
A new row appears in the dialog box.
3. In the new schedule row, in the **Retention** list, click a retention type.
4. If the schedule is for an Image Plug-in job that backs up volumes with SQL Server database files, do one of the following in the **SQL Application Settings** list:
 - To truncate database transaction logs after the backup, select **Truncate transaction logs**.
 - To run the backup without truncating logs, clear **Truncate transaction logs**.

Note: If you also back up databases using another tool (e.g., native SQL Server backup), use only one tool for truncating logs.

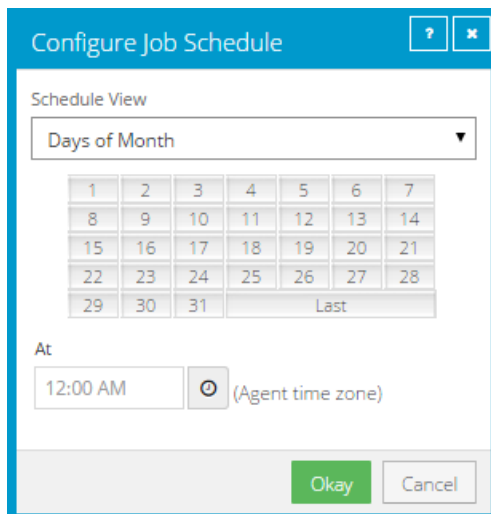
5. In the **Schedule** box, click the arrow.
The **Configure Job Schedule** dialog box opens.
6. In the **Configure Job Schedule** dialog box, do one of the following:

- To run the backup on specific days each week, in the **Schedule View** list, click **Days of Week**. Select the days when you want to run the job. Then use the **At** field to specify the time when you want to run the job.



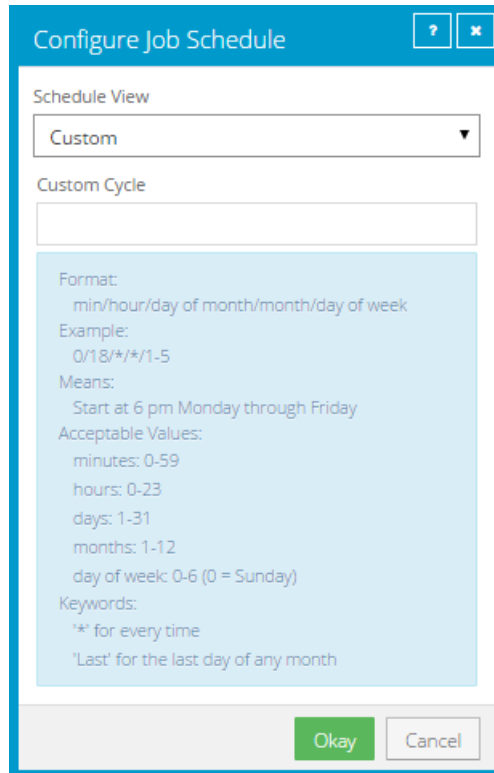
The screenshot shows the 'Configure Job Schedule' dialog box. The 'Schedule View' dropdown is set to 'Days of Week'. Below it, a row of buttons represents the days of the week: Su, Mo, Tu, We, Th, Fr, Sa. The buttons for Mo, Tu, We, Th, and Fr are highlighted with a blue background and a green checkmark, indicating they are selected. The 'At' field is set to '0:00 AM' and has a clock icon and '(Agent time zone)' next to it. At the bottom, there are 'Okay' and 'Cancel' buttons.

- To run the backup on specific dates each month, click **Days of Month** in the **Schedule View** list. On the calendar, select the dates when you want to run the job. Then use the **At** field to specify the time when you want to run the job.



The screenshot shows the 'Configure Job Schedule' dialog box. The 'Schedule View' dropdown is set to 'Days of Month'. Below it is a calendar grid with dates from 1 to 31, and a 'Last' option. The 'At' field is set to '12:00 AM' and has a clock icon and '(Agent time zone)' next to it. At the bottom, there are 'Okay' and 'Cancel' buttons.

- To create a custom schedule, click **Custom** in the **Schedule View** list. In the **Custom Cycle** dialog box, enter a custom schedule. Be sure to follow the format and notation as described.



7. Click **Okay**.

The new schedule appears in the **Schedule** box.

8. In the **Compression** list, click a compression level for the backup data. Compression levels optimize the volume of data sent against the speed of processing.
9. Do one of the following:
 - To allow the backup job to run without a time limit, click **None** in the Deferring list.
 - To specify a maximum amount of time that the backup job can run, click **Minutes** or **Hours** in the **Deferring** list. In the adjacent box, type the maximum number of minutes or hours that the job can run.

Note: When deferring is used, the backup job does not back up any new data after the specified amount of time, even if some data is not backed up. Changes to data that was previously backed up will be backed up, regardless of the amount of time specified.

10. To run the job on the specified schedule, select the **Enable** check box near the end of the row.
11. If there is more than one schedule row, you can use the **Priority** arrows to change the order of the schedule rows. Schedules higher in the list have a higher priority than schedules lower in the list.

If a job is scheduled to run at the same time by multiple schedules, the job only runs once at the scheduled time. If the schedules have different retention types, the job only runs with the retention type of the schedule that is highest in the list.

12. Click **Save**.

3.2 Run an ad-hoc backup

After a backup job is created, you can run the backup at any time, even if the job is scheduled to run at specific times.

To run an ad-hoc backup:

1. On the navigation bar, click **Computers**.

A grid lists available computers.

2. Find the computer with the backup job that you want to run, and expand its view by clicking the computer row.
3. Click the **Jobs** tab.
4. Find the job that you want to run, and click **Run Job** in its **Select Action** menu.

The **Run Job** dialog box shows the default settings for the backup.

Note: Beginning at this point, you can click **Start Backup** to immediately start the job. If you prefer, you can change backup options before running the job.

5. In the **Retention Scheme** list, click a retention type.

The retention type specifies the number of days a backup is kept on the vault, how many copies of a backup are stored online, and how long backup data is stored offline.

6. If you are backing up volumes with SQL Server database files using the Image Plug-in, do one of the following:
 - To truncate database transaction logs after the backup, select **Truncate transaction logs**.
 - To run the backup without truncating logs, clear **Truncate transaction logs**.

Note: If you also back up databases with another tool (e.g., native SQL Server backup), use only one tool for truncating logs.

7. Do one of the following:

- To allow the backup job to run without a time limit, clear the **Use Deferring** check box.
- To specify a maximum amount of time that the backup job can run, select the **Use Deferring** check box. From the **Backup time window** list, select **Minutes** or **Hours**. In the adjacent box, type the maximum number of minutes or hours that the job can run.

Note: When deferring is used, the backup job does not back up any new data after the specified amount of time, even if some data is not backed up. Changes to data that was previously backed up will be backed up, regardless of the backup time window.

Note: The **Use Deferring** check box is not available if you are backing up data to SSI (safeset image) files on disk.

8. Click **Start Backup**.

The **Process Details** dialog box shows the backup progress, and indicates when the backup is completed. Other recent job processes might also be listed in the dialog box. See [View current process information for a job](#).

9. If you want to stop the backup, click **Stop**.
10. To close the **Process Details** dialog box, click **Close**.

3.3 Synchronize a job

When a backup job is synchronized, the Agent checks which safesets for the job are online and available for restore.

A job is synchronized automatically when you restore data from the job. You can also synchronize a job manually at any time. A manual synchronization is recommended or required in the following cases:

- Before running backup jobs on reregistered computers.
- Before restoring data from jobs that are backed up to a Satellite vault and replicated to the cloud or another vault.
- To rebuild a delta (.dta) file for a job. If an error message in a log file says that the delta mapping file is corrupt, delete the delta (.dta) file from the job folder on the protected computer and then synchronize the job to rebuild the delta file.

To synchronize a job:

1. On the navigation bar, click **Computers**.
The Computers page shows registered computers.
1. Find the computer with the job that you want to synchronize. Expand its view by clicking its row.
2. Click the **Jobs** tab.
3. Find the job that you want to synchronize, and click **Synchronize** in its **Select Action** menu.

The **Process Details** dialog box shows the backup progress, and indicates when the backup is completed. Other recent job processes might also be listed in the dialog box. See [View current process information for a job](#).

4. If you want to stop the backup, click **Stop**.
To close the **Process Details** dialog box, click **Close**.

4 Restore from Image Plug-in backups

After backing up volumes from a Windows system using the Image Plug-in, you can restore:

- Entire volumes. See [Restore Windows volumes from an Image backup](#).
- Specific files and folders from the volumes. See [Restore files and folders from an Image backup](#).
- Systems from BMR backups, and volumes from BMR or non-BMR backups using the System Restore application. For more information, see the System Restore *User Guide*.

If you backed up SQL Server database files using Image Plug-in version 7.5 or later, you can restore database files from the application-consistent backups. See [Restore SQL Server database files or objects from Image backups](#).

Note: If you back up a volume that includes mount points, when you restore or mount the volume, the mount points will point to the original mount location.

4.1 Restore Windows volumes from an Image backup

After backing up volumes on a Windows computer using the Image Plug-in, you can restore volumes from the backup to selected live volumes (target volumes).

IMPORTANT: When you restore a volume from a backup, any data on the target volume will be lost.

A target volume must meet the following requirements:

- The target volume must be as large as or larger than the volume that was backed up.
- The Windows operating system cannot be installed on the target volume.
- The Windows Agent cannot be installed on the target volume.

After a restore, the system considers the target volume to be the size of the volume that was restored. For example, after you restore a 1 TB volume to a 2 TB target volume, the system sees the target as a 1 TB volume. Any remaining space on the target volume will not be usable.

To restore Windows volumes from an Image backup:

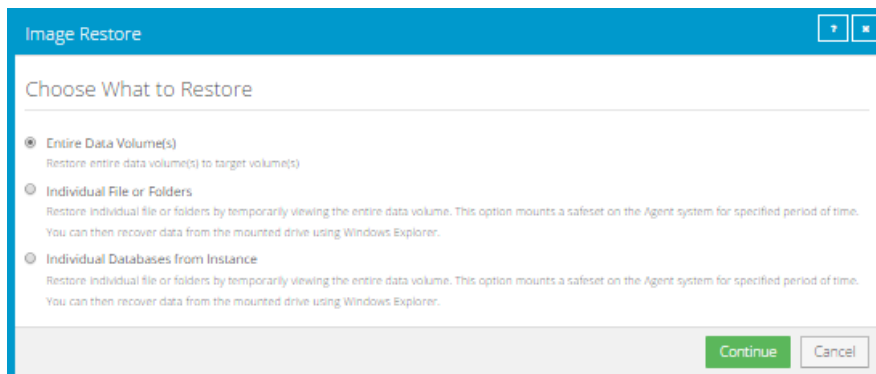
IMPORTANT: Before restoring volumes from an Image backup, stop any services on the system that are using the target volume (e.g., SQL Server or Exchange services). Restart the services after the restore is completed.

1. On the navigation bar, click **Computers**.
A grid lists available computers.
2. Find the Windows computer with the Image Plug-in, and expand its view by clicking the computer row.
3. Click the **Jobs** tab.

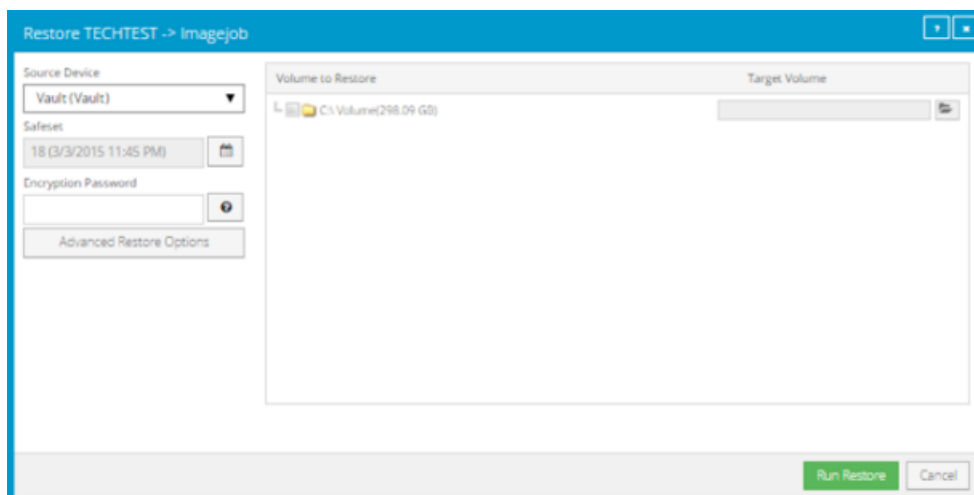
- Find the Image backup job with volumes that you want to restore, and click **Restore** in the job's **Select Action** menu.



Note: You can restore volumes from regular or Bare Metal Restore backups created using the Image Plug-in.

- In the **Image Restore** dialog box, select **Entire Data Volume(s)**, and then click **Continue**.




The **Restore** dialog box shows volumes that can be restored from the backup. The most recent safeset for the job appears in the **Safeset** box.




- To restore data from an older safeset, or from SSI (safeset image) files, do one of the following:
 - To restore data from an older safeset, click the calendar button.  In the calendar that appears, click the date of the safeset from which you want to restore. To the right of the calendar, click the specific safeset that you want to use.
 - To restore data from SSI (safeset image) files on disk, select **Directory on disk** from the **Source Device** list. Click the folder button.  In the **Select Folder** dialog box, select the directory where the files are located, and click **Okay**.

SSI files are full backups exported from the vault or backed up from a computer to disk instead of to a vault. It can be quicker to save backup files on physical media and transport them to a location for a restore, than to restore data from a vault in a remote datacenter.

Note: You cannot restore from backups to disk (SSI files) until the safeset is imported into the vault and the Agent is synchronized with the vault.

7. In the **Encryption Password** box, enter the encryption password. To view the password hint, click the **Hint** button. 
8. To change the log detail level or bandwidth options, click **Advanced Restore Options**. Specify settings in the **Advanced Restore Options** dialog box, and click **Okay**.
9. In the **Volume to Restore** column, select each volume that you want to restore.
10. For each selected volume, do the following to choose the live volume where it will be restored:

IMPORTANT: Data on the selected volume will be lost when the backed-up volume is restored.

- a. Click the folder icon. 

The **Select Volume** dialog box lists all live volumes on the computer. If you cannot restore the selected volume to a specific live volume (e.g., because the live volume is too small, contains the Windows operating system, or contains Windows Agent software), the volume is unavailable and cannot be clicked.

- b. Click the volume where you want to restore the backed up volume.
- c. Click **Okay**.

11. Click **Run Restore**.

The **Process Details** dialog box shows the restore progress and indicates when the restore is completed. The target volume goes offline until the backed-up volume is restored.

Other recent job processes might also be listed in the dialog box. See [View current process information for a job](#).

12. To close the **Process Details** dialog box, click **Close**. If the restore is running, it will continue to run.

4.2 Restore files and folders from an Image backup

After backing up volumes from a Windows server using the Image Plug-in, you can restore individual files and folders from the backup.

To restore files and folders from an Image backup, you mount volumes from a safeset as drives on the computer where you want to restore files and folders. You can then browse the drives using Windows Explorer, and copy files and folders that you want to restore.

Note: To restore files and folders from Image backups, Agent services must be running using the local system account.

Note: When SQL Server databases are backed up using Image Plug-in version 7.5 or later, you can also restore database files, tables and objects from the application-consistent backups.

To restore files and folders from an Image backup:

1. On the navigation bar, click **Computers**.

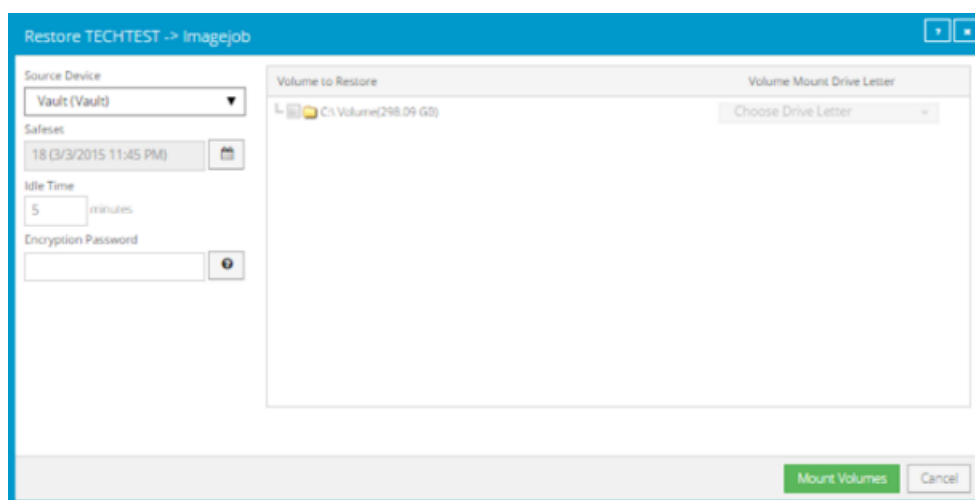
A grid lists available computers.



2. Find the Windows computer with the Image Plug-in, and expand its view by clicking the computer row.
3. Click the **Jobs** tab.
4. Find the Image backup job with files and folders to restore, and click **Restore** in the job's **Select Action** menu.

Note: You can restore files and folders from any Image Plug-in job.

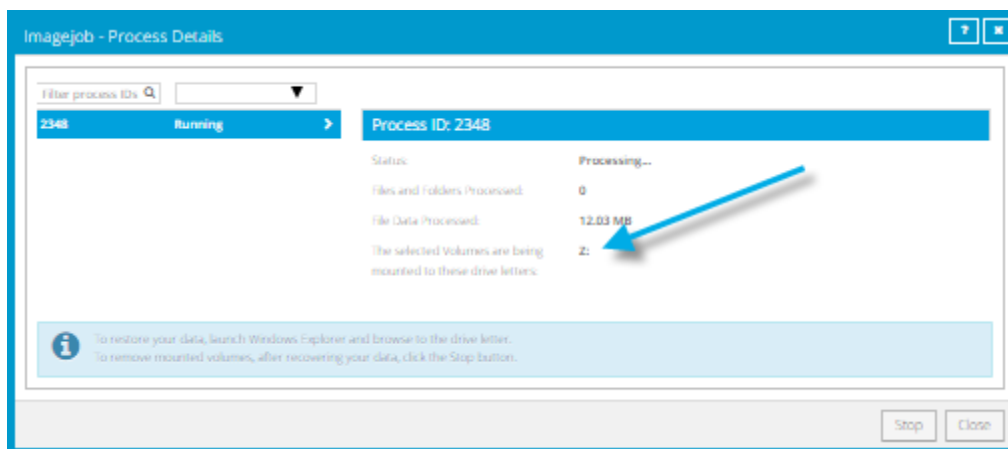
5. In the **Image Restore** dialog box, select **Individual File or Folders**, and then click **Continue**.

The **Restore** dialog box shows volumes that you can mount as drives. The most recent safeset for the job appears in the **Safeset** box.



6. To restore data from an older safeset, click the calendar button.  In the calendar that appears, click the date of the safeset from which you want to restore. To the right of the calendar, click the specific safeset that you want to use.
7. In the **Idle Time** box, enter the number of minutes of inactivity after which the shared drive will be unmounted automatically. The **Idle Time** can range from 2 to 180 minutes.
8. In the **Encryption Password** box, enter the encryption password. To view the password hint, click the **Hint** button. 
9. In the **Volume to Restore** column, select each backed-up volume from which you want to restore files and folders.
10. In the **Volume Mount Drive Letter** column, choose the drive letter for mounting each volume.
11. Click **Mount Volumes**.
12. If a confirmation message appears, read the message, and then click **Continue**.

The **Process Details** dialog box appears. When each volume is mounted, the drive letter is shown at the right side of the dialog box.



13. Use Windows Explorer to navigate to the drive or drives and copy files and folders that you want to restore.
14. When you are done restoring files and folders, click **Stop** to remove the mounted drives.

If you do not click **Stop**, the drive will be unmounted automatically after the number of minutes of inactivity specified in the **Idle Time** box. See Step [7](#).

4.2.1 Restore SQL Server database files or objects from Image backups



After backing up SQL Server databases using Image Plug-in version 7.5 or later, you can restore database files, tables and objects from the application-consistent backups.

To restore database files from Image Plug-in backups, you must mount a safeset as a drive on the computer where you want to restore database files. You can then restore files using Windows Explorer or tables and other objects using the Granular Restore for Microsoft Exchange and SQL application.

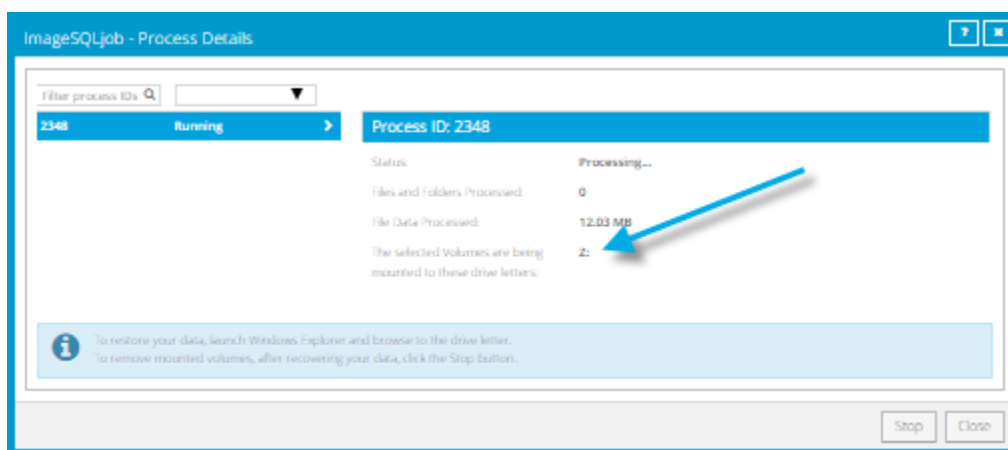
To restore SQL Server database files or objects from an Image Plug-in backup:

1. On the navigation bar, click **Computers**.
A grid lists available computers.
2. Find the computer with the SQL Server database backup created using the Image Plug-in, and expand its view by clicking the row for the computer.
3. Click the **Jobs** tab.
4. Find the job with the database you want to restore, and click **Restore** in the **Select Action** menu for the job.
The **Choose how to restore** dialog box appears.
5. Select **Individual File or Folders**.
6. Click **Continue**.

The **Restore** dialog box shows the most recent safeset for the job.

7. To restore data from an older safeset, click the calendar button.  In the calendar that appears, click the date of the safeset from which you want to restore. To the right of the calendar, click the specific safeset that you want to use.
8. In the **Idle Time** text box, enter the number of minutes of inactivity after which the share should automatically stop. This value can range from 2 to 180.
9. In the **Encryption Password** box, enter the data encryption password. To view the password hint, click the **Hint** button. 
10. In the **Volume to Restore** column, select each backed-up volume from which you want to restore files and folders.
11. In the **Volume Mount Drive Letter** column, choose the drive letter for mounting each volume.
12. Click **Mount Volumes**.
13. If a confirmation message appears, read the message, and then click **Continue**.

The **Process Details** dialog box appears. When each volume is mounted, the drive letter is shown at the right side of the dialog box.



14. Do one of the following:
 - To restore SQL Server databases, use Windows Explorer to browse to the location of the database files and copy the database files (.mdf and .ldf) that you want to restore. If a *databaseName_mod* folder exists, copy the *databaseName_mod* folder as well. You can then attach the database in Microsoft SQL Server.
 - To restore SQL Server database tables or objects, launch the Granular Restore for Microsoft Exchange and SQL application on a SQL Server system. In the Granular Restore application, choose **File > Open**, drill down into the volume to choose the .mdf file, then select and restore your data. For more information, see documentation for the Granular Restore for Microsoft Exchange and SQL application.
15. When you are done restoring database files, click **Stop** to remove the mounted drives.

If you do not click **Stop**, the drive will be unmounted automatically after the number of minutes of inactivity specified in the **Idle Time** box. See Step [8](#).

4.2.2 Restore databases in AlwaysOn Availability Groups

You must always restore a SQL Server database to the primary replica in an AlwaysOn Availability Group. If a Windows Agent and plug-in are not installed on the primary replica server, you must fail over to a server where the Agent and plug-in are installed before restoring the database.

After restoring a database to the primary replica and adding the database back into the AlwaysOn Availability Group, it will be replicated to the secondary replicas. To reduce the amount of replication traffic after a restore, you can run a “Restore from another computer” on any secondary replica server where the Windows Agent and plug-in are installed.

For information about backing up databases in AlwaysOn Availability Groups, see [Protect SQL Server databases in AlwaysOn Availability Groups](#).

To restore a primary database in an AlwaysOn Availability Group:

1. If the Agent and plug-in are not installed on the primary replica server, fail over to the secondary database instance where the Agent is installed.
The formerly secondary replica where you backed up the database becomes the primary replica.
2. Remove the primary database from the AlwaysOn Availability Group.
3. Delete the database from all secondary replicas.
4. Restore the primary database to the original database name using the Overwrite Existing Databases option.
5. Add the restored primary database to the AlwaysOn Availability Group using the Full Synchronization option.

After restoring a SQL Server database to the primary replica, to reduce the amount of required replication traffic, you can restore the database to secondary replica servers.

To restore a secondary database in an AlwaysOn Availability Group:

1. If you did not delete the database from all secondary replicas when restoring the primary database (see Step 3 in the previous procedure), remove the secondary database from the AlwaysOn Availability Group.
2. On a secondary replica server where the Agent and plug-in are installed, restore the database by running a Restore From Another Computer using the No Recovery option.
3. Add the restored secondary database to the AlwaysOn Availability Group using the Join option.

5 Monitor computers and processes

You can monitor backups, restores, and protected computers using the following Portal features:

- **Computer page.** The Computer page shows status information for protected computers and their jobs. See [View computer and job status information](#).
- **Process Details dialog box.** This dialog box shows information about all running, queued and recently-completed processes for a job. See [View current process information for a job](#).
- **Process logs and safeset information.** Process logs indicate whether each backup and restore completed successfully, and provide information about any problems that occurred. You can also view information about the safeset created by a specific backup. See [View a jobs process logs and safeset information](#).
- **Monitor page.** The Monitor page shows the most recent backup status for each job, and allows you to navigate to the computer and job for each backup. See [View and export recent backup statuses](#).

5.1 View computer and job status information

On the Computer page in Portal, you can view status information for protected computers and their jobs.





To view computer and job status information:


1. On the navigation bar, click **Computers**.

The Computers page shows registered computers.

The **Availability** column indicates whether each computer is online or offline. Online computers are in contact with Portal, while offline computers are not currently available. A computer can be offline if it is turned off, if the Agent has been uninstalled from the system, or if the system has been lost.

The **Status** column shows the status of each computer. Possible statuses include:









-  OK — Indicates that all jobs on the computer ran without errors or warnings.
 -  Warning — Indicates that one or more of the computer's jobs completed with warnings.
 -  Attention — Indicates that one or more of the computer's jobs failed or completed with errors.
 -  Unconfigured — Indicates that no jobs have been created for the computer.
2. Find the computer for which you want to view logs, and click the row to expand its view.
 3. View the **Jobs** tab.

If a backup or restore is running for a job, an “In Progress” symbol  appears beside the job name, along with the number of processes that are running.

	Name	Job Type	Description
1	AppAware	Image	
2	FilesAndFolders	Local System	

If you click the symbol, the **Process Details** dialog box shows information about running, queued and recently-completed processes for the job. See [View current process information for a job](#).

The **Last Backup Status** column shows the result of the last backup attempt for each job. Possible statuses include:

-  Completed — Indicates that the last backup completed successfully, and a safeset was created.
-  Completed with warnings — Indicates that the last backup completed and a safeset was created, but problems occurred during the backup. For example, a warning could indicate that a file or volume that was selected in the backup job was not available for backup.
-  Deferred — Indicates that the last backup was deferred. A safeset was created, but not all data that was selected was backed up.
- Deferring is used to prevent large backups from running at peak network times. When deferring is enabled, a backup job does not back up any new data after a specified amount of time.
-  Missed — Indicates that the job has not run for 7 days.
-  Completed with errors — Indicates that the backup completed and a safeset is available for restore, but problems occurred. Typically, this status indicates that not all of the data was backed up.
-  Failed — Indicates that the backup failed and no safeset was created.
-  Never Run — Indicates that the backup job has never run.
-  Cancelled


To view logs for a job, click the job status. For more information, see [View a jobs process logs and safeset information](#).

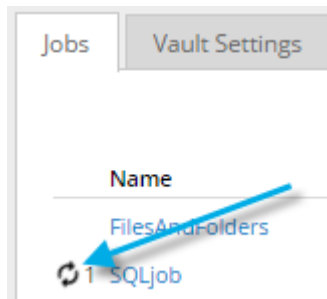
5.2 View current process information for a job

In the Process Details dialog box, you can view information about running, queued and recently-completed processes for a job. Processes include backups, restores and synchronizations. Process information is typically deleted within an hour after the process ends.


To view current process information for a job:

1. Do one of the following:
 - On the Computers page, on the Jobs tab, start a backup, restore or synchronization.

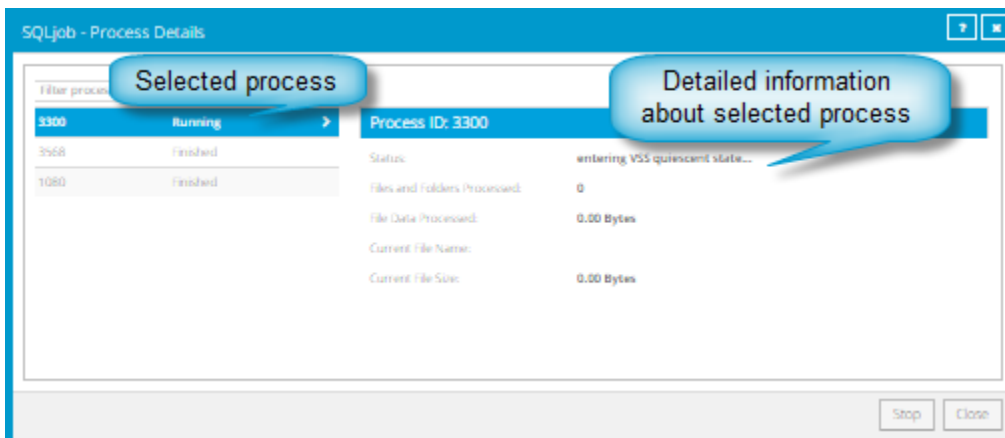
- On the Computers page, on the Jobs tab, click the “In Progress” symbol  beside the job name.



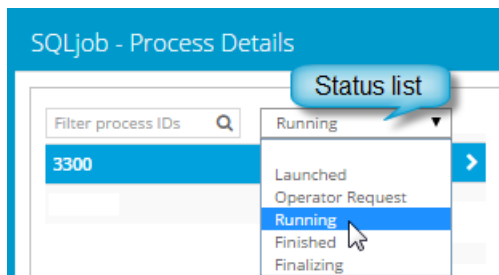
- On the Monitor page, click the “In Progress” symbol  beside the job name.



The **Process Details** dialog box lists processes that are running, queued and recently completed for the job. Detailed information is shown for the process that is selected on the left side of the dialog box.



- To view information about a different process, click the process on the left side of the dialog box. Detailed information for the process is shown at the right side of the dialog box.
- To show only some processes in the dialog box, do one of the following in the status list:
 - To only show queued processes, click **Launched**.
 - To only show processes that are waiting for user action, click **Operator Request**.
 - To only show processes that are in progress, click **Running**.
 - To only show completed processes, click **Finished**.
 - To only show processes that are finishing, click **Finalizing**.



5.3 View a job's process logs and safeset information

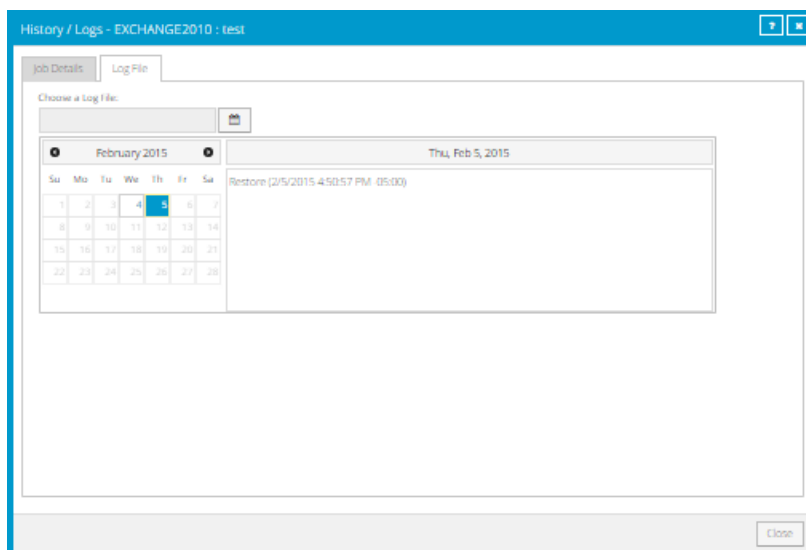
To determine whether a backup or restore completed successfully, or to determine why a process failed, you can view a job's process logs.


You can also view information about safesets created for the job. A safeset is an instance of backup data on the vault. For most Agents, one safeset is created by each successful backup.

To view a job's process logs and safeset information:

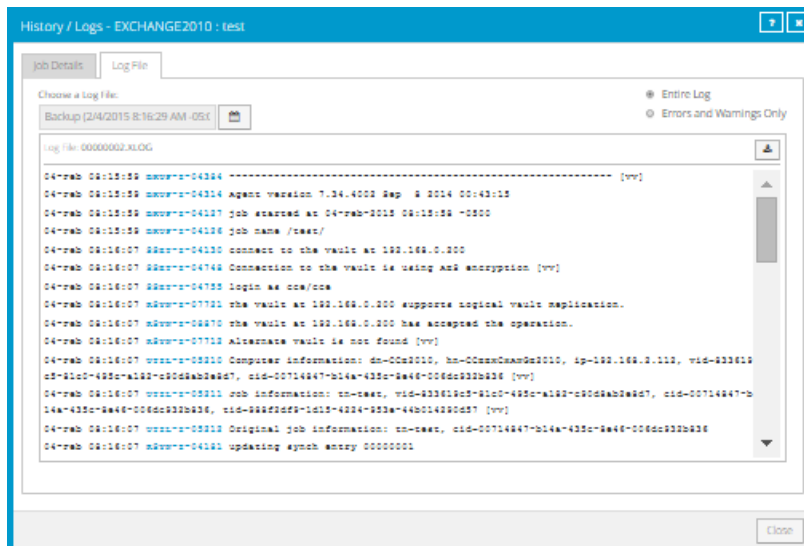
1. On the navigation bar, click **Computers**.
The Computers page shows registered computers.
2. Find the computer for which you want to view logs, and click the row to expand its view.
On the **Jobs** tab, the **Last Backup Status** column shows the status of each backup job.
3. To view log files for a job, do one of the following:
 - In the job's **Select Action** menu, click **History / Logs**.
 - In the **Last Backup Status** column, click the job status.

The **History / Logs** window lists the most recent backups, restores and synchronizations on the computer.




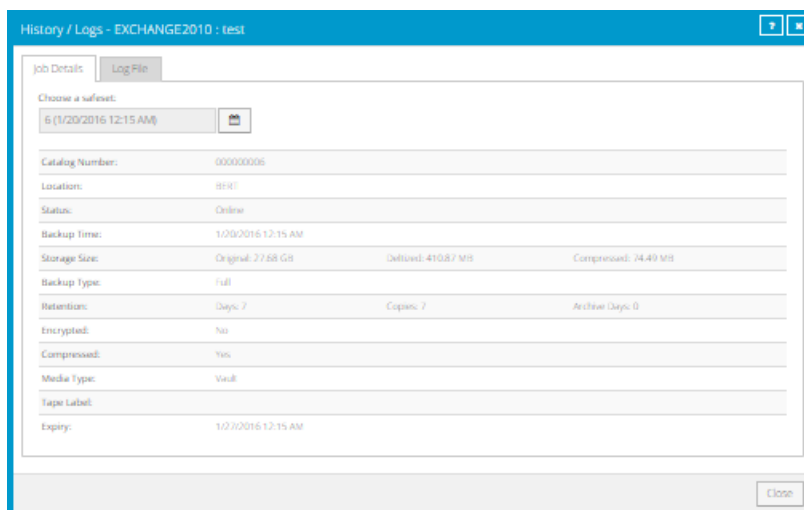
- To view processes for a different day, click the calendar button.  In the calendar that appears, click the date of the log that you want to view. In the list of processes on the selected date, click the process for which you want to view the log.

The **History / Logs** window shows the selected log.



- To only show errors and warnings in the log, click the **Errors and Warnings Only** option at the top right of the window.
- To view safeset information for a particular backup, click the **Job Details** tab. The tab shows safeset information for the job's most recent backup.

To view information for a different safeset, click the calendar button.  In the calendar that appears, click the date of the backup for which you want to view information. In the list of backups on the selected date, click the backup for which you want to view information. The tab shows safeset information for the selected backup.



5.4 View and export recent backup statuses

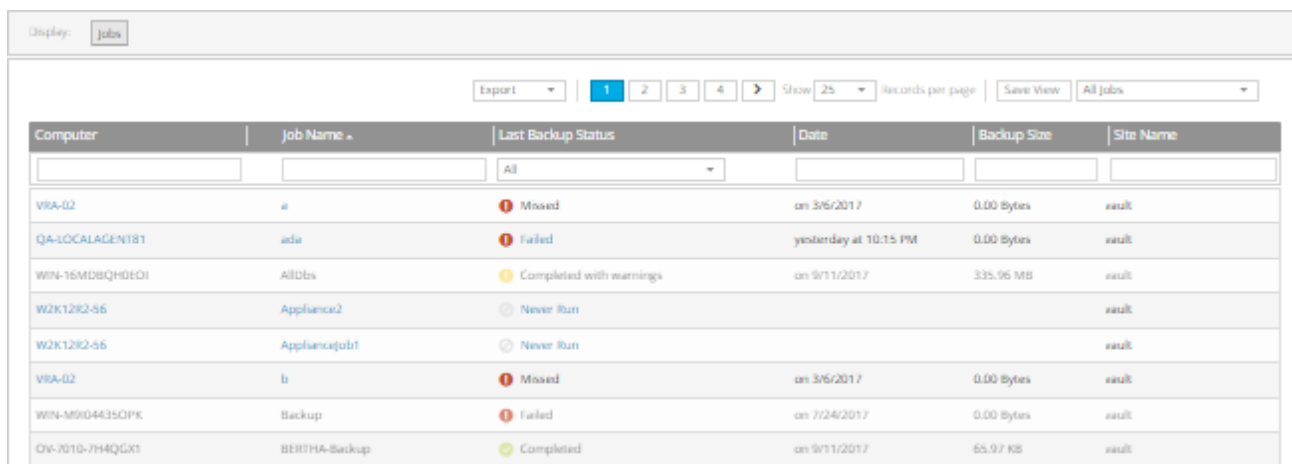
You can view recent backup statuses for computers on the Monitor page in Portal. You can also export the information in comma-separated values (.csv), Microsoft Excel (.xls), or Adobe Acrobat (.pdf) format.

From the Monitor page, you can navigate to related information on the Computers page or in the Logs window.

To view and export recent backup statuses:


1. On the navigation bar, click **Monitor**.

The Monitor page shows recent backup statuses for jobs in your site.



Computer	Job Name	Last Backup Status	Date	Backup Size	Site Name
VBA-02	a	Missed	on 3/6/2017	0.00 Bytes	naul
QA-LOCALAGENT81	ada	Failed	yesterday at 10:15 PM	0.00 Bytes	naul
WIN-16MD8QH0E01	AllDts	Completed with warnings	on 9/11/2017	335.06 MB	naul
W2K12R2-56	Appliance2	Never Run			naul
W2K12R2-56	Appliancejob1	Never Run			naul
VBA-02	b	Missed	on 3/6/2017	0.00 Bytes	naul
WIN-M9D4435DPK	Backup	Failed	on 7/24/2017	0.00 Bytes	naul
OV-1010-7H4QG01	BERTHA-Backup	Completed	on 9/11/2017	65.97 KB	naul

2. To change which backup statuses appear on the page, click the views list at the top of the page, and then click the view that you want to apply.
3. To view information for a job or computer on the Computers page, click the name of an online computer or job.
4. To view the job's logs in the History/Logs window, click the job's last backup status.
5. To export backup status information from the page, click the **Export** box. In the list that appears, click one of the following formats for the exported data file:
 - CSV (comma-separated values)
 - XLS (Microsoft Excel)
 - PDF (Adobe Acrobat)



The data file is downloaded to your computer in the specified format.

6 Carbonite Server Backup Support

If you have a question about Carbonite Server Backup that isn't covered in this guide, our frequently-updated Knowledge Base contains comprehensive information. The Knowledge Base is your first stop when searching for any Carbonite Server Backup solutions you may need. We highly recommend searching here first for the quickest answers to your questions.

Knowledge Base: <http://support.carbonite.com/evault>

What can we help you with?


Search

Popular Searches
[pending reboot](#), [restore](#), [clnt-e-04103](#)

6.1 Contacting Carbonite

If you need live assistance from a qualified support agent, Carbonite Support is here for you 24 hours a day, 7 days a week (excluding US holidays). Please feel free to get in touch with us, and we'll help out any way we can! You can find the contact information for Carbonite Support in the Knowledge Base:

<http://support.carbonite.com/evault>



Tip: When contacting Support with a technical issue, please have both the program's log files and the store you are having difficulty with ready.

To gather log files, click **File** menu and choose *Open log folder*. Compress the contents of the folder in a .zip file and attach it to your support request.

If the log archive and/or mail store exceeds 10MB, you may not be able to send them as an email attachment. In that case, upload instructions will be provided to you upon request.