EVault Software

SQL Server Plug-in 8.3

User Guide

Revision: This manual has been updated for Version 8.3.

Software Version: 8.30 (November 2016)

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 note d.

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

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.

The EVault Software Agent, EVault Software CentralControl, and EVault Software 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 EVault Software Agents and EVault Software Director applications also have the added security feature of an over the wire encryption method.

Contents

1 Introduction to the SQL Server Plug-in

To protect Microsoft SQL Server databases, install the SQL Server Plug-in with the Windows Agent on the machine where SQL Server is running. You can then add and run backup jobs that specify which SQL Server databases to back up and where to save the backup data.

The SQL Server Plug-in can back up databases that span volumes, databases that have Transparent Data Encryption (TDE) enabled and databases in AlwaysOn Availability Groups. The Plug-in can also back up BLOB data from filestream-enabled databases. You can run full database backups, full database with transaction logs backups or transaction log only backups.

> *Note:* You can only back up transaction logs for databases that use the full or bulk-logged recovery model.

When installed with the Cluster Support Plug-in, the SQL Server Plug-in can protect databases on SQL Server clusters. For more information about the Cluster Support Plug-in, see the Portal help or Windows Agent guide.

After a SQL Server database is backed up, you can restore the database to the SQL Server instance where it was backed up, to a different instance or to flat files on disk. In a single pass, the Plug-in can restore a database from a full backup and transaction log backups to its state at a selected point in time.

You can also use the SQL Server Plug-in to back up Microsoft SharePoint 2013 and 2010 databases. You can restore entire SharePoint databases or restore individual SharePoint items (e.g., site collections, web sites, lists, documents) using the SQL Server Plug-in and the Granular Restore for Microsoft SharePoint application.

For installation and configuration information, see the Windows Agent guide or Portal online help. For supported platform information, see the Windows Agent release notes.

Note: You can also back up SQL Server databases using the Windows Agent and Image Plug-in version 7.5 or later. For more information, see the Image Plug-in guide or Portal online help.

1.1 Required permissions

In addition to permissions required for the Windows Agent, the account specified during the Agent and SQL Server Plug-in installation must have the public server role to perform full SQL Server backups.

The account must have the "sysadmin" role to perform transaction log backups.

2 Add a SQL Server Plug-in backup job

After a Windows computer with the SQL Server Plug-in is added and configured in Portal, you can create a backup job for one or more databases in a SQL Server instance. The backup job specifies which database or databases to back up, and where to save the backup data. A SQL Server Plug-in job cannot include databases from multiple SQL Server instances.

You can also back up a SharePoint 2013 or 2010 database with a SQL Server Plug-in job.

When you create a SQL Server database backup job, you must specify Windows administrator or SQL Server administrator credentials that allow the Agent to connect to the instance where you are backing up databases.

To back up the data, you can run the backup job manually or schedule the job to run. When scheduling or running a job, you can specify whether to back up the database, the transaction logs, or both. See Run and schedule backups and synchronizations.

From the backup, you can restore an entire database. You can also use a Granular Restore application to restore specific items from the database. See Restore items from a SQL Server or SharePoint database.

To add a SQL Server database backup job:

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

   The Computers page shows registered computers.

2. Find a Windows computer with the SQL Server Plug-in, and expand its view by clicking the computer row.

3. Click the **Jobs** tab.

   If a valid vault connection is not available for the computer, you cannot access the **Jobs** tab. For information about adding a vault connection, see the Portal help.

4. In the **Select Job Task** menu, click **Create New SQL Server Job**.

5. In the **Connect to SQL Server** dialog box, specify the following information:

   - In the **Instance** list, select the SQL Server instance where you want to back up databases.
   - To connect to the instance using a Windows administrator account, select **Windows authentication**.
   - To connect to the instance using a SQL Server administrator account, select **SQL authentication**.
   - In the **User Name** box, type the user name for connecting to the instance.
   - In the **Password** box, type the password of the specified user.
   - If you selected Windows authentication, in the **Domain** box, type the domain of the specified account.

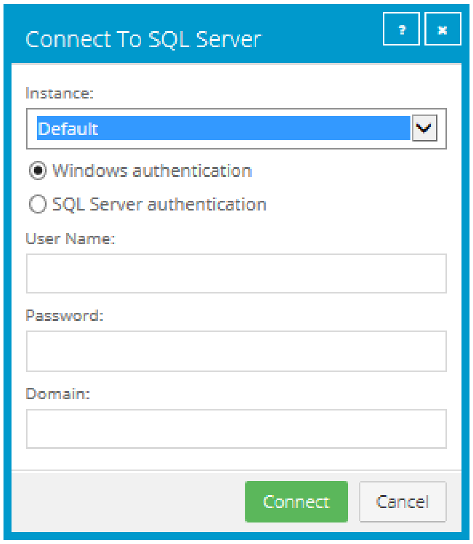

Instance: Default

● Windows authentication

○ SQL Server authentication

6. Click **Connect**.
7. 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. For more information, see Log file options.
   - 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.

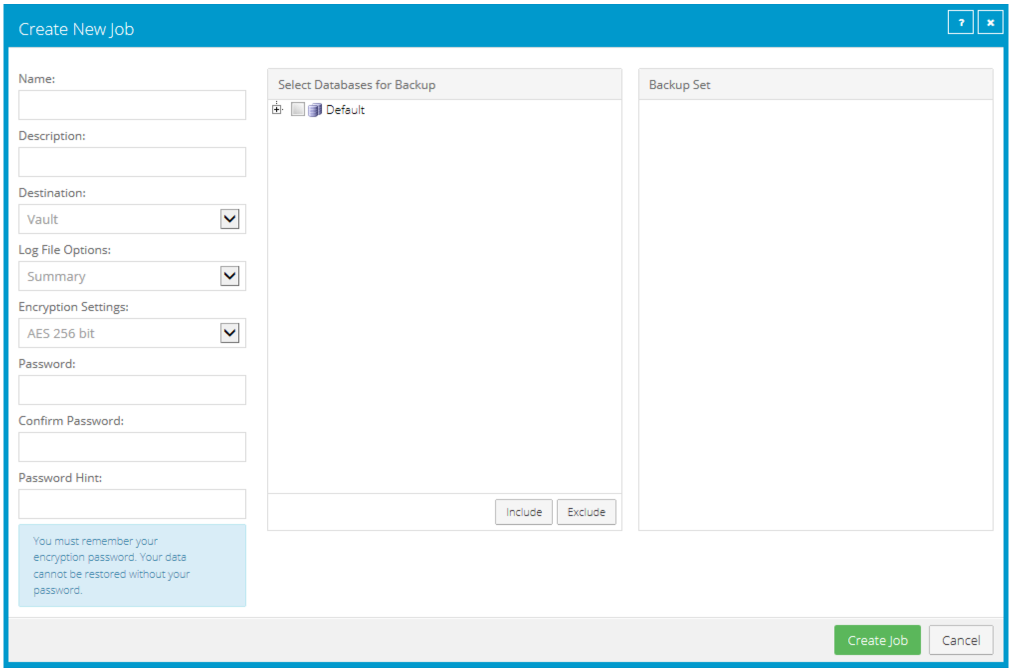

Destination: Vault

Log File Options: Summary

Encryption Settings: AES 256 bit

You must remember your encryption password. Your data cannot be restored without your password.

8. In the **Select Databases for Backup** box, do one or more of the following to add databases to the backup job:

   - To add specific databases to the backup job, select the check box for each database, and then click **Include**. The included databases appear in the **Backup Set** box.

   - To back up all databases in the selected SQL Server instance, select the check box for the instance, and then click **Include**. The included instances appear in the **Backup Set** box.

     Note: When the job runs, newly-added databases in the selected instance are automatically backed up.

   - To back up databases with names that match a filter when the job runs, select the check box for the SQL Server instance, and then click **Include**. An inclusion record with an asterisk (*) appears in the **Backup Set** box.

     In the **Database Filter** box, enter the names of databases to include. Separate multiple names with commas, and use asterisks (*) and question marks (?) as wildcard characters. For example, to back up databases with names that end with “Management” or include the word “database” followed by a single character, enter the following filter: *management, database?

     Note: Filters are applied when the backup job runs. New databases that match the specified filters are automatically backed up when the job runs.

     Note: Filters are not case-sensitive.

9. To exclude databases from the backup job, do one or more of the following in the **Select Databases for Backup** box:

   - To exclude specific databases from the backup job, select the check box for each database, and then click **Exclude**. The excluded databases appear in the **Backup Set** box.

- To exclude databases with names that match a filter when the backup job runs, select the check box for the SQL Server instance, and then click **Exclude**. A record with an asterisk (*) appears in the **Backup Set** box.

  In the **Database Filter** box, enter the names of databases to exclude. Separate multiple names with commas, and use asterisks (*) and question marks (?) as wildcard characters. For example, to exclude databases if their names begin with "M", enter the following filter: m*

  Note: Filters are applied when the backup job runs. New databases that match the specified filters are automatically excluded when the backup job runs.

  Note: Filters are not case-sensitive.

10. To remove an inclusion or exclusion record from the **Backup Set** box, click the Delete button beside the record. 🗑
11. Click **Apply Now** to consolidate and simplify records in the **Backup Set** box, if changes need to be applied.
12. 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. See Run and schedule backups and synchronizations.

    Click **Cancel** if you do not want to create a schedule at this time.

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.

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. 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.
- 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. 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.

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. 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 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.

For information about restoring SQL Server databases in AlwaysOn Availability Groups, see Restore databases in AlwaysOn Availability Groups.

2.2 Protect SQL Server clusters

To protect a SQL Server cluster, you must install the Windows Agent with the Cluster Support Plug-in and SQL Server Plug-in on each node in the cluster. In Portal, you can then register a virtual server for the SQL Server role in Portal and create and run backup jobs on the virtual server. Backup jobs on a virtual server are automatically directed to the active cluster node and will not reseed after a failover.

To fully protect a SQL Server cluster, you must back up:

- the quorum disk
- each physical node in the cluster
- cluster volumes
- the SQL Server databases to provide point-in-time database recovery.

When a cluster is fully protected, you can recover the cluster if components are lost, are corrupted or fail.

For detailed information, see Windows cluster information in the Portal online help or the Windows Agent guide.

2.3 Log file options

When you create or edit a backup job, you can specify the level of detail for job logging. Select one of the following job logging levels from the list:

- **Files** — Provides the most detailed information, and is typically used for troubleshooting. Provides information about files that are backed up.

- **Directory** — Provides less detail than the Files logging level. Provides information about folders that are backed up.
- **Summary** — Provides high-level information, including the vault and Agent version, and backup sizes.
- **Minimal** — Provides high-level information, including the vault and Agent version.

Changing the logging level only affects log files that are created from that point and after. It does not affect previously-created log files.

2.4 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 a SQL Server Plug-in backup job is deferred while a database is being backed up, the backup for that database is incomplete and the database cannot be restored. However, you can restore databases that were completely backed up in the job before the job was deferred. If the backup is deferred while a SharePoint database is being backed up, the backup is incomplete and you cannot restore items from the database.

  Note: Backups to SSI files on disk cannot be deferred.
- For a SQL Server Plug-in backup job, you can specify whether to back up the database, the transaction logs, or both. Frequent transaction log backups are recommended for databases with a high level of activity.

  Note: After a transaction log backup, logs are marked for truncation. If you also back up databases using another tool (e.g., native SQL Server backup), use only one tool for truncating logs.

  Note: Transaction logs can only be backed up for databases that use the full or bulk-logged recovery model.

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.

When scheduling multiple SQL Server database jobs in the same instance, it is good practice to schedule the jobs so that their running times do not overlap. Simultaneous backups are supported, but are not recommended.

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. In particular, try to avoid overlapping schedules for SQL Server database jobs in the same instance. Simultaneous backups in the same SQL Server instance are supported, but are not recommended.

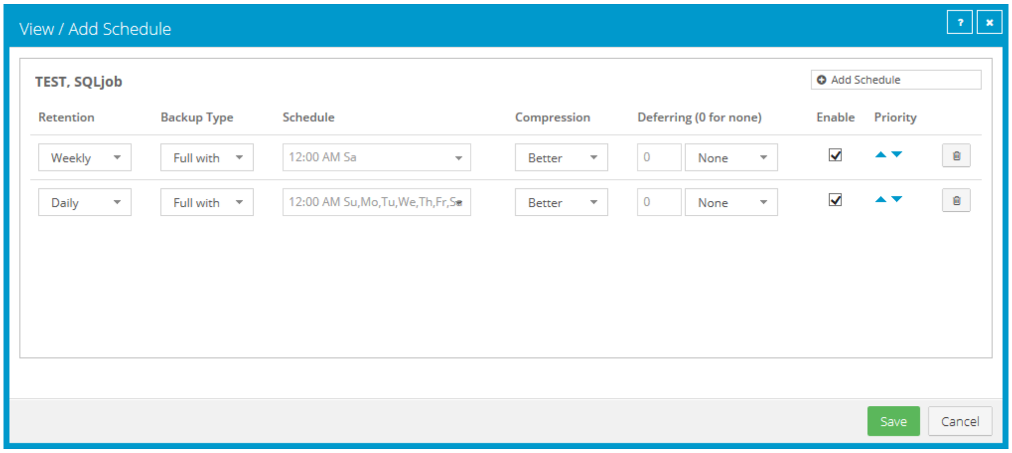

TEST, SQLjob

Retention	Backup Type	Schedule	Compression	Deferring (0 for none)		Enable	Priority
Weekly	Full with	12:00 AM Sa	Better	0	None	☑	▲▼
Daily	Full with	12:00 AM Su,Mo,Tu,We,Th,Fr,Sa	Better	0	None	☑	▲▼

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. Do one of the following in the **Backup Type** list:

   - To back up each database from the point in time when the backup starts, click **Full**.
   - To back up each database and its transaction logs from the point in time when the backup starts, click **Full with transaction logs**.
   - To back up the database transaction logs only from the point in time when the backup starts, click **Transaction logs only**. When **Transaction Logs only** is selected, the entire database and its transaction logs will be backed up when the job first runs. In subsequent backups, only the transaction logs will be backed up.

   Note: After a transaction log backup, logs are marked for truncation. If you also back up databases using another tool (e.g., native SQL Server backup), use only one tool for truncating logs.

   Note: Transaction logs can only be backed up for databases that use the full or bulk-logged recovery model.

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.

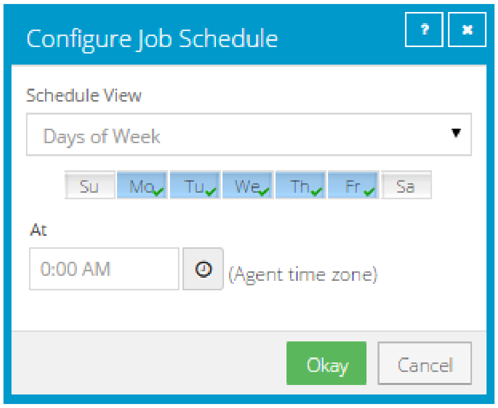

Schedule View: Days of Week

Su | Mo ✓ | Tu ✓ | We ✓ | Th ✓ | Fr ✓ | Sa

At: 0:00 AM (Agent time zone)

   - 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.

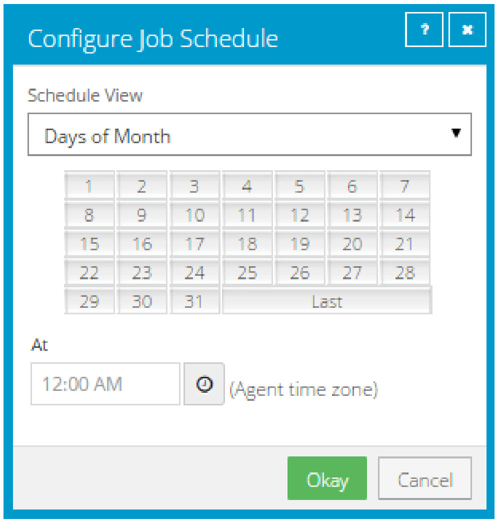

Schedule View: Days of Month

1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30	31	Last			

At: 12:00 AM (Agent time zone)

- 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.

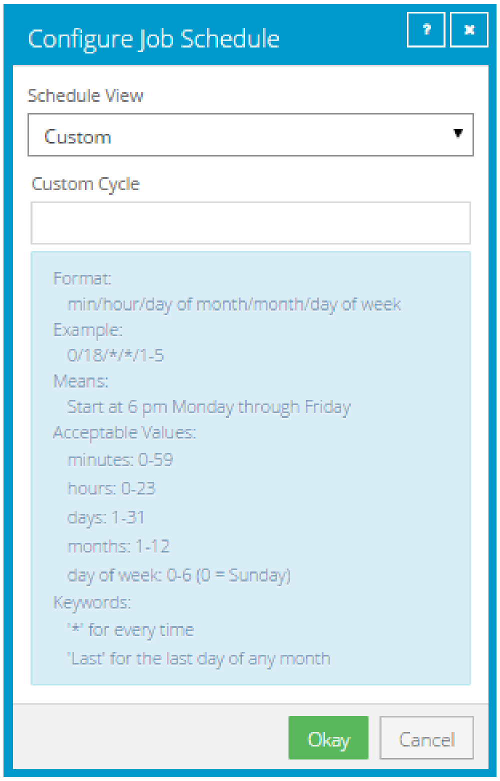

Schedule View: Custom

Custom Cycle

Format:
min/hour/day of month/month/day of week
Example:
0/18/*/*/1-5
Means:
Start at 6 pm Monday through Friday
Acceptable Values:
minutes: 0-59
hours: 0-23
days: 1-31
months: 1-12
day of week: 0-6 (0 = Sunday)
Keywords:
'*' for every time
'Last' for the last day of any month

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.

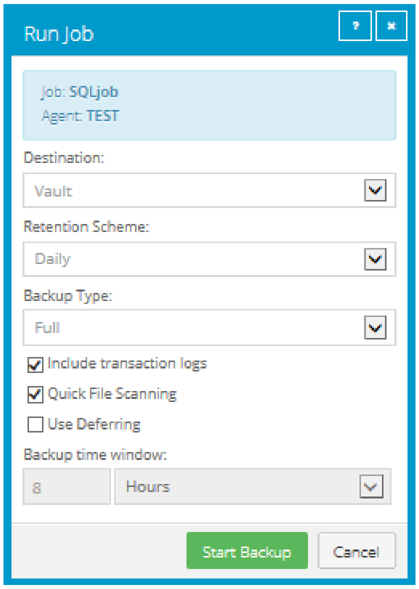
Run Job

Job: SQLjob
Agent: TEST

Destination: Vault

Retention Scheme: Daily

Backup Type: Full

☑ Include transaction logs
☑ Quick File Scanning
☐ Use Deferring

Backup time window: 8 Hours

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. Do one of the following:

   - To back up the database, click **Full**. To also back up the database's transaction logs, select **Include transaction logs**.

   - To back up transaction logs only, click **Transaction Log**. When **Transaction Log** is selected, the database and its transaction logs will be backed up when the backup first runs. In subsequent backups, only the transaction logs will be backed up.

     Note: After a transaction log backup, logs are marked for truncation. If you also back up databases using another tool (e.g., native SQL Server backup), use only one tool for truncating logs.

     Note: Transaction logs can only be backed up for databases that use the full or bulk-logged recovery model.

7. To enable Quick File Scanning, select the **Quick File Scanning** check box.

   Quick File Scanning (QFS) reduces the amount of data read during the backup process. Any file streams that have not changed since the last backup are skipped. Without QFS, files are read in their entirety. Note that changes in delta-file format might cause QFS to be temporarily disabled during the first backup following an upgrade. This could cause this first backup to take longer than usual.

8. 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.

9. 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.

10. If you want to stop the backup, click **Stop**.

11. 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.

2. Find the computer with the job that you want to synchronize. Expand its view by clicking its row.

3. Click the **Jobs** tab.

4. 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.

5. If you want to stop the backup, click **Stop**.

   To close the **Process Details** dialog box, click **Close**.

4 Restore SQL Server databases

After backing up SQL Server databases using the SQL Server Plug-in, you can restore databases directly to a SQL Server instance, or restore databases to flat files. See Restore databases directly to SQL Server or Restore SQL Server databases to flat files.

When restoring a SQL Server database in an Always On Availability Group, you must always restore the database to the primary replica. See Restore databases in AlwaysOn Availability Groups.

4.1 Restore databases directly to SQL Server

After backing up SQL Server databases using the SQL Server Plug-in, you can restore databases directly to a SQL Server instance.

If transaction logs have been backed up using an alternative method (e.g., native SQL Server backup), you can restore a database in the restoring state so that you can apply transaction logs to the database after the restore.

When restoring system databases, the **master** database must be restored first, by itself. Other system databases can then be restored.

You must specify a Windows or SQL Server administrator account for connecting to SQL Server during a restore.

After restoring a SQL Server 2016 database that is stretched to Microsoft Azure, you must run a stored procedure (sys.sp_rda_reauthorize_db) to reconnect the local restored database to the remote Azure data. See “Restore the connection between the SQL Server database and the remote Azure database” on the Microsoft Developer Network website: https://msdn.microsoft.com/en-us/library/mt733205.aspx#reconnect

To restore a database directly to SQL Server:

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

   A grid lists available computers.

2. Find the computer with the SQL Server database backup that you want to restore, and expand its view by clicking the row for the computer.
3. Click the **Jobs** tab.
4. Find the job with the database that you want to restore. In the job’s **Select Action** menu, click **Restore**.
5. In the **Choose how to restore** dialog box, select **Restore database to a SQL Server instance**.
6. In the **Instance** list, click the SQL Server instance where you want to restore the database.
7. Do one of the following:

- To connect to the instance using a Windows administrator account, select **Windows authentication**. Enter the user name, password, and domain in the appropriate fields.
- To connect to the instance using a SQL Server administrator account, select **SQL Server authentication**. Enter the user name and password in the appropriate fields.

8. Click **Continue**.

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

9. 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.

10. In the **Database Selection** box, select the check box for each database that you want to restore.

11. In the **Encryption Password** box, enter the data encryption password. To view the password hint, click the **Hint** button.

12. Do one of the following:

    - To restore one or more databases with their original names, select **Original Database Names**.
    - To restore one database with a new name, select **Alternate Database Name**. In the field that appears, enter the new name for the restored database.

      Note: You can only restore one database if **Alternate Database Name** is selected.

13. Do one of the following:

    - To overwrite the existing database if you restore a database with the same name as the existing database, select **Overwrite existing databases**.
    - To fail the restore if a database with the same name already exists, clear **Overwrite existing databases**.

      If **Overwrite existing databases** is not selected, and you are restoring multiple databases, the restore fails for all databases if even one database has the same name as an existing database.

14. To restore the database in restoring state, select **Restore using No Recovery option**.

If this option is selected, and transaction logs have been backed up using an alternative method (e.g., native SQL Server backup), you can apply transaction logs to the database after it has been restored.

15. To specify an alternate location for database files, select **Alternate Path**. Click the folder button. 📅 In the **Select Folder** dialog box, select the alternate file location, and click **Okay**.

    Note: The alternate file location is only used if the original location for database files is not available.

16. To change the log detail level, bandwidth throttling setting or hard recovery option, click **Advanced Restore Options**. In the dialog box, do one or more of the following:

    - In the **Log Level Detail** list, select the level of detail for job logging.
    - Select or clear the **Use all available bandwidth** option.

17. Click **Run Restore**.

    The **Process Details** dialog box shows the restore progress and indicates when the restore is completed. Other recent job processes might also be listed in the dialog box. See View current process information for a job.

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

4.2 Restore SQL Server databases to flat files

After backing up SQL Server databases using the SQL Server Plug-in, you can restore a SQL Server database to flat files. SQL Server tools can then be used to bring the data into a database.

To restore a SQL Server database to flat files:

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

    A grid lists available computers.

2. Find the computer with the SQL Server database backup that you want to restore, 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.
5. In the **Choose how to restore** dialog box, select **Restore to folder**.
6. Click **Continue**.

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

7. 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.

8. In the **Database Selection** box, select the check box for each database that you want to restore.
9. In the **Encryption Password** box, enter the data encryption password. To view the password hint, click the **Hint** button. ❓
10. Under **Restore Destination**, enter a path for the destination, or click the folder button. 📂 In the **Select Folder** dialog box, select the location where you want to restore, and click **Okay**.
11. To change the log detail level, bandwidth throttling setting or hard recovery option, click **Advanced Restore Options**. In the dialog box, do one or more of the following:
    - In the **Log Level Detail** list, select the level of detail for job logging.
    - Select or clear the **Use all available bandwidth** option.
12. Click **Run Restore**.

    The **Process Details** dialog box shows the restore progress and indicates when the restore is completed. Other recent job processes might also be listed in the dialog box. See View current process information for a job.

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

4.3 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 Restore items from a SQL Server or SharePoint database

If a Microsoft SharePoint 2010 or 2013 database is backed up using the SQL Server Plug-in, you can restore items such as site collections, websites, lists and documents from the backup.

If a Microsoft SQL Server database is backed up using the SQL Server Plug-in or Image Plug-in, you can restore specific tables and objects from the backup.

To restore items from a database backup, you must first use Portal to expose the safeset as a shared resource. You can then use a Granular Restore application to find and restore items from the backup. To restore items from a SharePoint database backup, use Granular Restore for Microsoft SharePoint. To restore items from a SQL Server database backup, use Granular Restore for Microsoft Exchange and SQL. For more information, or to obtain a Granular Restore application, contact your service provider.

To restore items from a SQL Server or SharePoint database:

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

   A grid lists available computers.

2. Find the computer with the safeset with SharePoint or SQL Server data that you want to restore, and expand its view by clicking the computer row.
3. Click the **Jobs** tab.
4. Find the job with the SharePoint data that 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 **Restore items to a SharePoint or SQL Server database**, and click **Continue**.

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

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 **Encryption Password** box, enter the data encryption password. To view the password hint, click the **Hint** button.
8. In the **Idle Time** box, enter the number of minutes of inactivity after which the share should automatically stop. The value can range from 2 to 180 minutes.
9. Select or clear the **Use all available bandwidth** option.
10. Click **Share**.

    The **Process Details** dialog box shows the status of the share process. When the share is available, the share path appears at the right side of the dialog box.

11. Click the Copy Path to Clipboard button. The path is now available for you to paste into the Granular Restore application.
12. Do one of the following:
    - To restore SharePoint items, launch the Granular Restore for Microsoft SharePoint application on a SharePoint 2010 or 2013 system.
    - To restore SQL Server database items, launch the Granular Restore for Microsoft Exchange and SQL application on a SQL Server system.
13. Paste the path for the SQL safeset share into the Granular Restore application.
14. Select and restore your data. For more information, see documentation for the Granular Restore application.
15. When you no longer need to share the safeset, click **Stop**.

    When you click **Stop** or the share idle time is reached, the **Process Details** dialog box indicates that the share is no longer available.

6 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.

6.1 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 s.

   The **Availability** column indicates whether each 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 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.

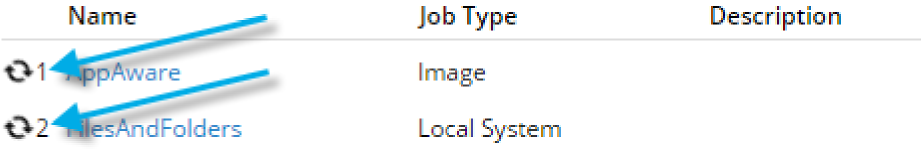

Name	Job Type	Description
↻1 AppAware	Image	
↻2 FilesAndFolders	Local System	

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](#).

6.2 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.

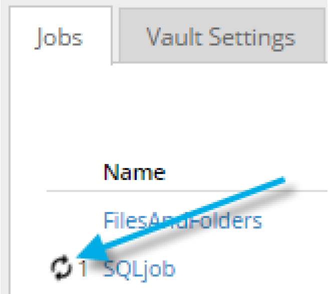

Name
FilesAndFolders
⟳ 1 SQLjob

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

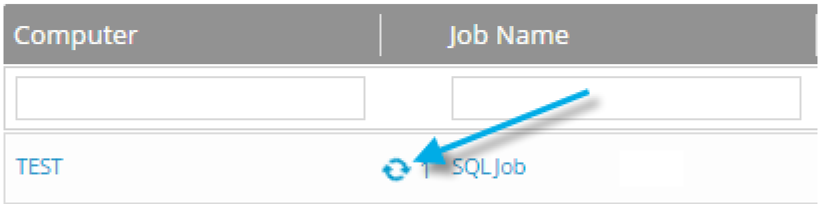

Computer	Job Name
TEST	⟳ 1 SQL Job

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.

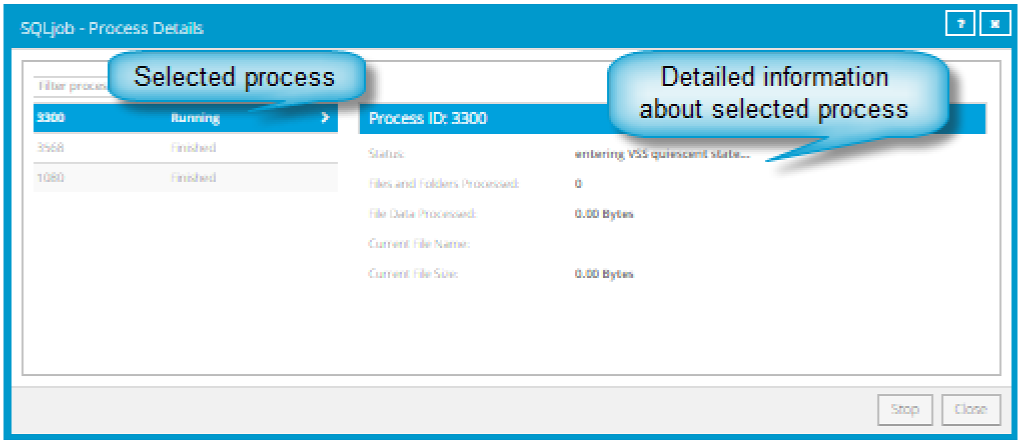

3300	Running
3568	Finished
1080	Finished

Process ID: 3300

Status:	entering VSS quiescent state...
Files and Folders Processed:	0
File Data Processed:	0.00 Bytes
Current File Name:	
Current File Size:	0.00 Bytes

2. 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.

3. 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**.

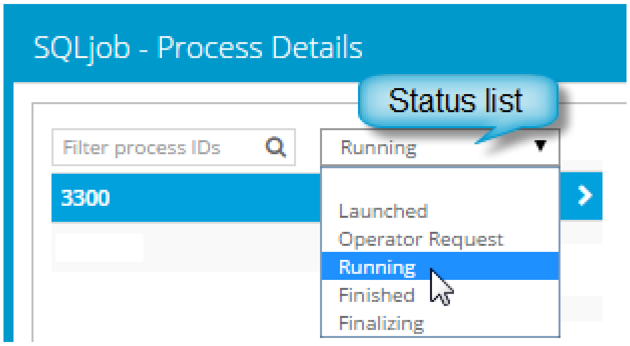

6.3 View a job’s process logs and safeset information

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 s.

2. Find the 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.

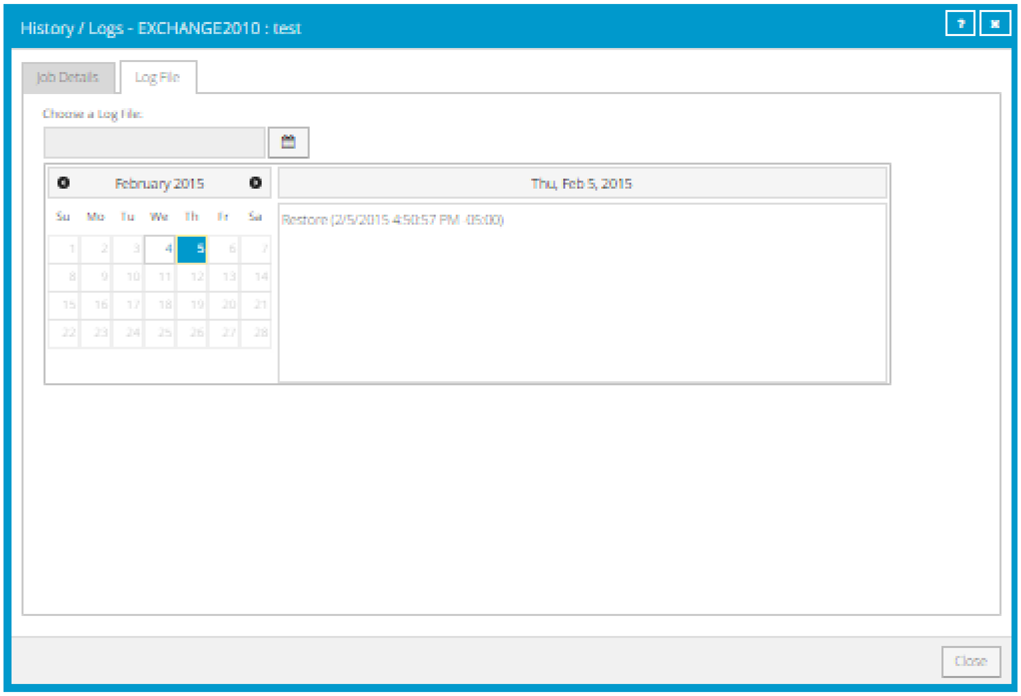

Su	Mo	Tu	We	Th	Fr	Sa
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28

Restore (2/5/2015 4:50:57 PM -05:00)

4. 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.

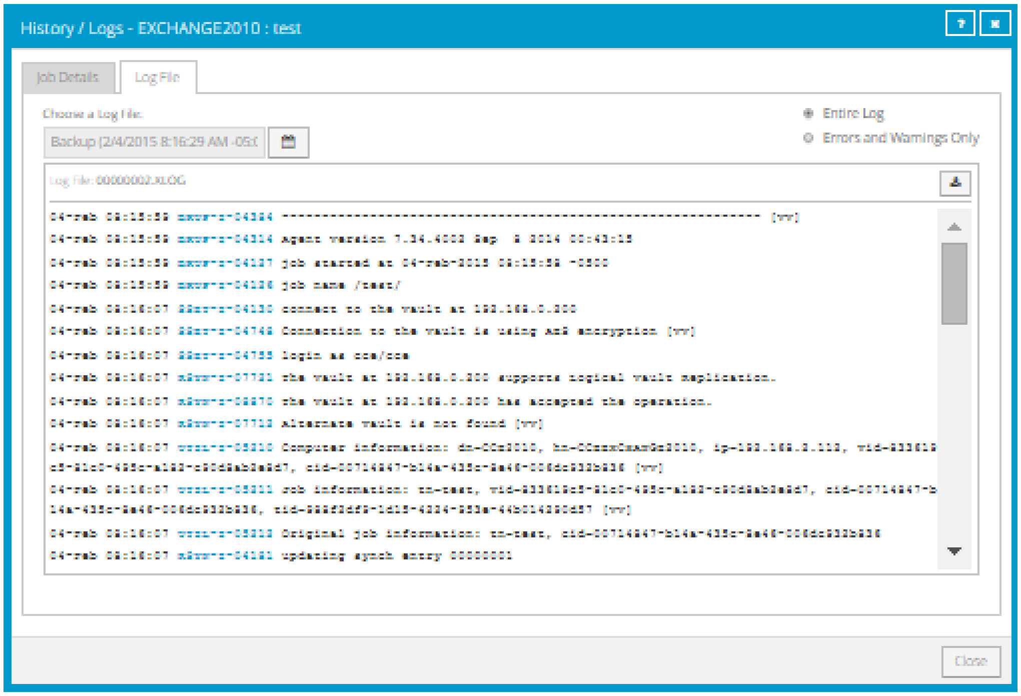

● Entire Log
○ Errors and Warnings Only

Backup (2/4/2015 8:16:29 AM -05:0

Log File: 00000002.XLOG

```
04-Feb 08:15:58 BKUP-I-04384 ------------------------------------------------------------ [vv]
04-Feb 08:15:58 BKUP-I-04314 agent version 7.34.4002 Sep  8 2014 00:43:15
04-Feb 08:15:58 BKUP-I-04127 job started at 04-Feb-2015 08:15:58 -0500
04-Feb 08:15:58 BKUP-I-04128 job name /test/
04-Feb 08:16:07 SSET-I-04130 connect to the vault at 192.168.0.200
04-Feb 08:16:07 SSET-I-04748 Connection to the vault is using AES encryption [vv]
04-Feb 08:16:07 SSET-I-04755 login as cca/cca
04-Feb 08:16:07 SVTR-I-07721 the vault at 192.168.0.200 supports logical vault replication.
04-Feb 08:16:07 SVTR-I-08870 the vault at 192.168.0.200 has accepted the operation.
04-Feb 08:16:07 SVTR-I-07712 alternate vault is not found [vv]
04-Feb 08:16:07 VSSX-I-05310 Computer information: dn-CCE2010, hn-CCEXCHANGE2010, ip-192.168.2.112, vid-833619
c5-81c0-485c-a182-c80d8ab2e8d7, cid-00714847-b14a-435c-8a48-008dc832b836 [vv]
04-Feb 08:16:07 VSSX-I-05311 Job information: tn-test, vid-833619c5-81c0-485c-a182-c80d8ab2e8d7, cid-00714847-b
14a-435c-8a48-008dc832b836, tid-888f2df8-1d15-4224-853a-44b014280d57 [vv]
04-Feb 08:16:07 VSSX-I-05312 Original job information: tn-test, cid-00714847-b14a-435c-8a48-008dc832b836
04-Feb 08:16:07 BKUP-I-04181 updating synch entry 00000001
```

5. To only show errors and warnings in the log, click the **Errors and Warnings Only** option at the top right of the window.

6. 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.

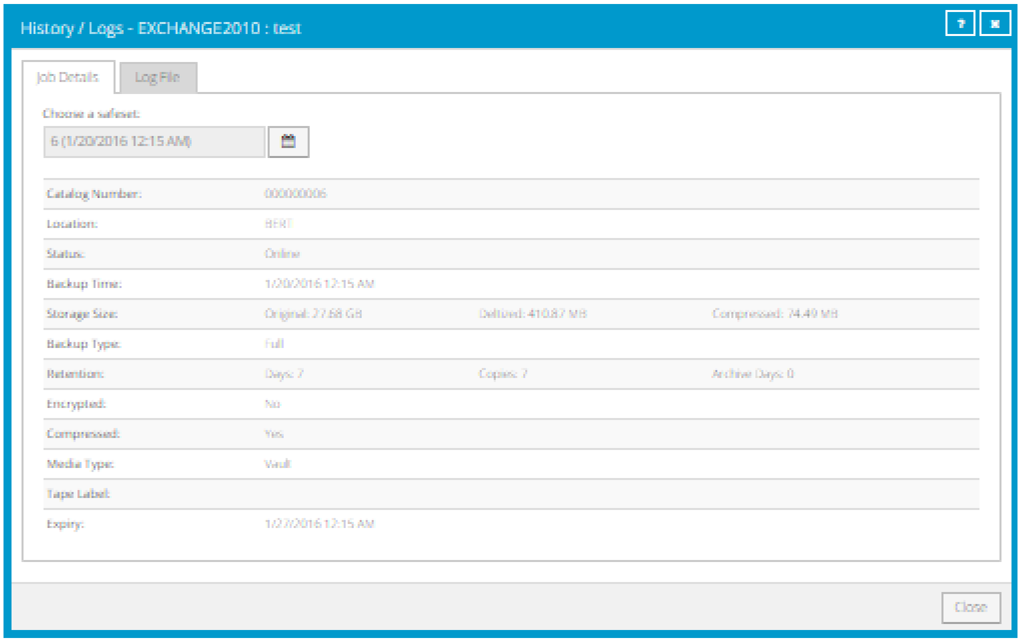

Choose a safeset: 6 (1/20/2016 12:15 AM)

Catalog Number:	000000006		
Location:	BERT		
Status:	Online		
Backup Time:	1/20/2016 12:15 AM		
Storage Size:	Original: 27.68 GB	Deltized: 410.87 MB	Compressed: 74.49 MB
Backup Type:	Full		
Retention:	Days: 7	Copies: 7	Archive Days: 0
Encrypted:	No		
Compressed:	Yes		
Media Type:	Vault		
Tape Label:			
Expiry:	1/27/2016 12:15 AM		

6.4 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.

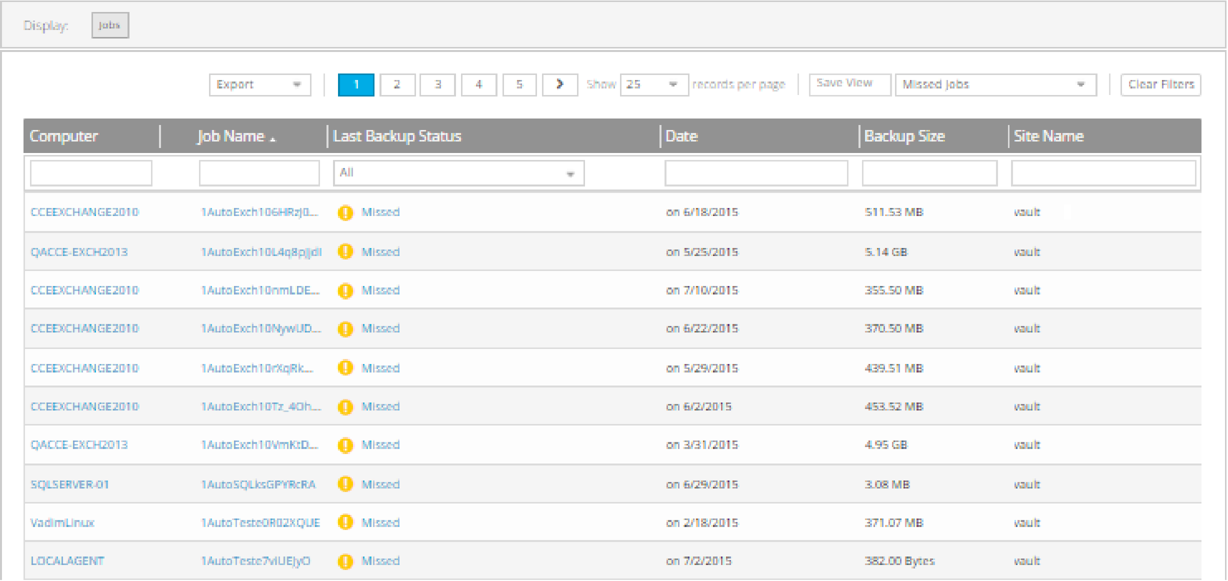

Computer	Job Name	Last Backup Status	Date	Backup Size	Site Name
		All			
CCEEXCHANGE2010	1AutoExch106HRzj0...	Missed	on 6/18/2015	511.53 MB	vault
QACCE-EXCH2013	1AutoExch10L4q8pjjdl	Missed	on 5/25/2015	5.14 GB	vault
CCEEXCHANGE2010	1AutoExch10nmLDE...	Missed	on 7/10/2015	355.50 MB	vault
CCEEXCHANGE2010	1AutoExch10NywUD...	Missed	on 6/22/2015	370.50 MB	vault
CCEEXCHANGE2010	1AutoExch10rXqRk...	Missed	on 5/29/2015	439.51 MB	vault
CCEEXCHANGE2010	1AutoExch10Tz_4Oh...	Missed	on 6/2/2015	453.52 MB	vault
QACCE-EXCH2013	1AutoExch10VmKtD...	Missed	on 3/31/2015	4.95 GB	vault
SQLSERVER-01	1AutoSQLksGPYRcRA	Missed	on 6/29/2015	3.08 MB	vault
VadimLinux	1AutoTeste0R02XQUE	Missed	on 2/18/2015	371.07 MB	vault
LOCALAGENT	1AutoTeste7viUEjyO	Missed	on 7/2/2015	382.00 Bytes	vault

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)

   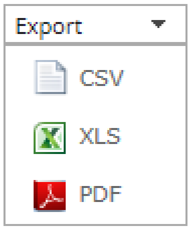

   The data file is downloaded to your computer in the specified format.