

Exchange Plug-in 8.7

User Guide



OSP·EVAULT

© Copyright 2019. All Rights Reserved.

The software manufacturer 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, the software manufacturer reserves the right to revise this publication and to make changes from time to time in the content hereof without obligation of the software manufacturer 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.

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 (C) 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 Agent, Agent Console, and Vault applications have the added encryption option of 128/256 bit 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). See: <http://csrc.nist.gov/encryption/aes/round2/r2report.pdf> for details.

The Agent and Vault applications have the added security feature of an over the wire encryption method.

Document History

Version	Date	Description
1	December 2018	Initial guide provided for Exchange Plug-in 8.7x.
2	March 2019	Changed references to "Exchange 2010, 2013 and 2016" to "Exchange 2010 and later versions". Beginning with Exchange Plug-in 8.71, Exchange 2019 is also supported with the Exchange Plug-in.

Contents

1	Introduction to the Exchange Plug-in	2
1.1	Required permissions.....	2
2	Add an Exchange backup job.....	3
3	Run and schedule backups and synchronizations.....	6
3.1	Schedule a backup.....	7
3.2	Specify whether scheduled backups retry after a failure	10
3.3	Run an ad-hoc backup.....	11
3.4	Plan Full and Incremental Exchange backups	12
3.5	Synchronize a job	13
4	Restore Exchange databases.....	14
4.1	Restore databases in Exchange 2010 and later versions	15
4.2	Restore Exchange 2007 databases.....	16
4.3	Restore Exchange databases to flat files.....	16
5	Restore Exchange mailboxes, messages and other objects.....	18
6	Monitor computers, jobs and processes.....	20
6.1	View computer and job status information	20
6.2	View an unconfigured computer's logs.....	22
6.3	View current process information for a job	23
6.4	Monitor backups using email notifications	25
6.5	View a job's process logs and safeset information	27
6.6	View and export recent backup statuses	30

1 Introduction to the Exchange Plug-in

To protect Microsoft Exchange databases, install the Exchange Plug-in with the 64-bit Windows Agent on the server where Exchange is running. You can then add and run backup jobs that specify which Exchange databases to back up, and where to save the backup data.

After an Exchange database is backed up, you can:

- Restore the database to its original location, to an alternate Exchange database location or to flat files.
- Restore mailboxes, messages, and other Exchange objects using the plug-in and the Granular Restore for Microsoft Exchange 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.

1.1 Required permissions

In addition to permissions required for the Windows Agent, the account specified during the Agent and Exchange Plug-in installation must belong to the following groups:

- Exchange Organization Administrators
- Group Policy Owners
- Schema Admins
- Enterprise Admins
- Domain Admins

2 Add an Exchange backup job

After a Windows computer with the Exchange Plug-in is added in Portal, you can create a backup job for one or more Microsoft Exchange databases. The backup job specifies which databases (in Exchange 2010 and later Exchange versions) or storage groups (in Exchange 2007) to back up, and where to save the backup data.

When running or scheduling an Exchange backup job, you can specify whether to run a Full or Incremental backup and whether to validate the Exchange data. See [Run and schedule backups and synchronizations](#) and [Plan Full and Incremental Exchange backups](#).

After an Exchange backup job runs successfully, transaction logs for databases in the job are truncated so that the logs only contain changes that occurred after the backup.

When an Exchange server has multiple databases (Exchange 2010 and later versions) or storage groups (Exchange 2007), you can put the databases and storage groups into separate jobs and run the jobs simultaneously. Do not create parallel jobs for the same database (Exchange 2010 and later versions) or storage group (Exchange 2007) or conflicts could prevent the jobs from completing successfully. Conflicts could also occur if you create backups using third-party applications or Agents on other Database Availability Group members.

With Exchange 2007, your high-availability solution can affect the backup strategy. When Cluster Continuous Replication (CCR) is used, a backup only succeeds if there is only one database in the storage group in the backup job. You can run backups on either the Active or Passive node, but running backups on the Passive node is recommended. When Local Continuous Replication (LCR) is used, you can specify whether to back up data on the active copy of each storage group or on the replica copy. The replica copy can only be used for a backup if LCR is enabled for all storage groups in the job. If LCR is not enabled for all storage groups, the active copy is used for the backup.

When an Exchange backup job runs, databases in the job that are mounted or healthy are backed up. Other databases are skipped. If a database is skipped when a job runs but is mounted or healthy for the following run, the database backup does not reseed during the following run. However, if a database is skipped in two or more consecutive runs, the database backup reseeds during the next backup when the database is mounted or healthy. If no databases in a backup job are mounted or healthy when the job runs, the backup fails.

To add an Exchange backup job:

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

The Computers page shows registered computers.

2. Find a Windows computer with the Exchange 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 more information, see the Portal online help or Windows Agent guide.

4. In the **Select Job Task** menu, click **Create New Exchange Server 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.
- For new backup jobs, the encryption method is AES 256 bit. Existing jobs can have other encryption methods.
- In the **Password** and **Confirm Password** boxes, enter an encryption password. You can also enter a password hint in the **Password Hint** box.

6. In the **Include in Backup** box, do one of the following:

- To add specific Exchange databases or storage groups to the backup job, select the check box for each database or storage group, and then click **Include**.
- To back up Exchange databases or storage groups that match a filter when the job runs, select the check box for the server, and then click **Include**. An inclusion record appears in the **Backup Set** box.

In the **Filter** box, enter the names of databases or storage groups to back up. 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: Available items depend on the Microsoft Exchange version. You can select databases for Exchange 2010 and later versions. You can select storage groups for Exchange 2007.

Note: Filters in a backup job are applied when the job runs. New databases or storage groups that match the filters are automatically backed up when the job runs.

7. To back up data on an Exchange 2007 server where local continuous replication (LCR) is enabled, do one of the following:
 - To only back up data from the active copy of each storage group, select **Only back up active instance**.
 - To back up data from the replica copy of each storage group, clear **Only back up active instance**.

Note: The replica copy can only be used for a backup if LCR is enabled for all storage groups in the job. If LCR is not enabled for all selected storage groups, the active copy of each storage group is used for the backup.

8. 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](#).

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.

Note: Incremental backups for Exchange cannot be deferred, even if deferring is enabled. Deferring can be applied to full backups for Exchange.

- For an Exchange Plug-in backup job, you can specify whether to:
 - Run a Full or Incremental backup. When the Full backup type is selected, the database files, checkpoint file and transaction logs are backed up. When the Incremental backup type is selected, the database files, checkpoint file and transaction logs are backed up in the first “seed” backup, but only the checkpoint file and transaction logs are backed up in subsequent runs. For more information, see [Plan Full and Incremental Exchange backups](#).
 - Validate Exchange data during the backup. When this option is selected, a utility checks the Exchange data during the backup. If data corruption is detected, the backup fails and the corruption is reported.

For computers with Windows or Linux Agent version 8.60 or later, you can specify whether scheduled backups should automatically retry if they do not run successfully. You can also specify how many times a scheduled backup should retry after a failed attempt, and specify the amount of time between retries. See [Specify whether scheduled backups retry after a failure](#).

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**.

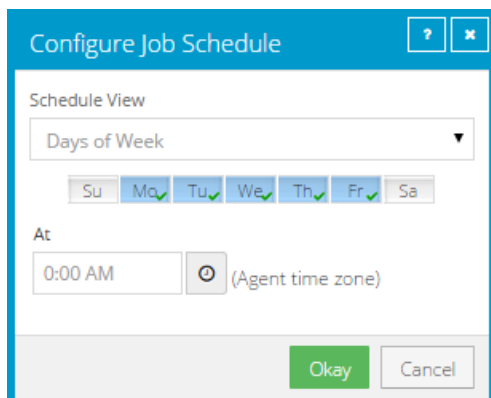
A new row appears in the dialog box.
 3. In the new schedule row, in the **Retention** list, click a retention type.
 4. If the schedule is for an Exchange database backup job, do the following:
 - In the **Backup Type** list, do one of the following:
 - To only back up transaction logs and the checkpoint file after the first “seed” backup, click **Incremental**.
 - To back up the database files, checkpoint file and transaction logs, click **Full**.
- For more information, see [Plan Full and Incremental Exchange backups](#).
- To validate Exchange data during the backup, select **Validate Exchange database**.

- In the **Schedule** box, click the arrow.

The **Configure Job Schedule** dialog box opens.

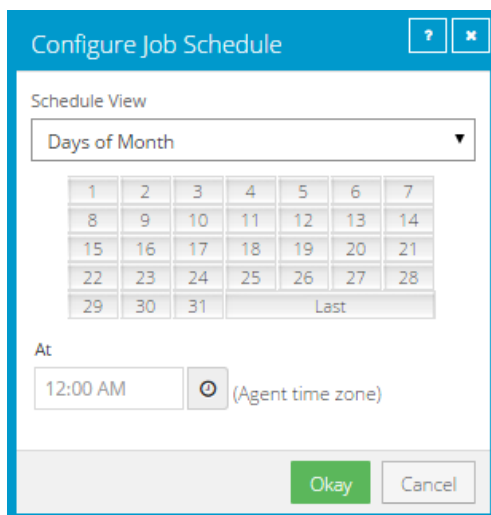
- In the **Configure Job Schedule** dialog box, do one of the following:

- To run the backup on specific days each week, select **Days of Week** in the **Schedule View** list. 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 in blue with 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)' text next to it. At the bottom, there are 'Okay' and 'Cancel' buttons.

- To run the backup on specific dates each month, select **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)' text next to it. At the bottom, there are 'Okay' and 'Cancel' buttons.

- To create a custom schedule, select **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.

Configure Job Schedule

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

Okay Cancel

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.

Note: Incremental backups for Exchange cannot be deferred, even if deferring is enabled. Deferring can be applied to full backups for Exchange.

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.

- If an Automatic Retry for Scheduled Backups section appears at the bottom of the View / Add Schedule dialog box, you can specify whether scheduled backups should retry after a failed backup. See [Specify whether scheduled backups retry after a failure](#).
- Click **Save**.

3.2 Specify whether scheduled backups retry after a failure

You can specify whether scheduled backups automatically retry if they do not run successfully.

You can also specify how many times a scheduled backup should retry after a failed attempt, and specify the amount of time between retries.

Note: Automatic retry settings only apply to scheduled backups. A backup will not retry automatically after a failed ad-hoc backup attempt.

To specify whether scheduled backups retry after a failure:

- Do one of the following:
 - On the navigation bar, click **Computers**. Find the computer for specifying automatic retry settings, and click the computer row to expand its view. On the **Jobs** tab, in the **Select Action** menu for a job, click **View/Add Schedule**.
 - Create a new backup job. The **View/Add Schedule** dialog box appears when you save the job.
- In the Automatic Retry for Scheduled Backups section, do one of the following:
 - To specify that scheduled backups should not retry after failed backup attempts, clear the **Retry failed job** check box.
 - To specify that scheduled backups should retry after failed backup attempts, select the **Retry failed job** check box. In the **Number of retries** box, enter the number of times the backup should try again. In the **Wait before each retry attempt for [] minutes** box, enter the number of minutes that the Agent should wait before the next backup attempt.

The screenshot shows the 'View / Add Schedule' dialog box for a backup job named 'PROTECTEDSERVER, CloudServerBackup'. The dialog has a blue header and a white body. At the top right, there is an 'Add Schedule' button. Below the header, there are several sections for configuring the backup job. The 'Automatic Retry for Scheduled Backups' section is highlighted with a red box. It contains a checked checkbox for 'Retry failed backup', a 'Number of retries' field set to 1, and a 'Wait before each retry attempt for' field set to 1 minute. At the bottom right of the dialog, there are 'Save' and 'Cancel' buttons.

- Click **Save**.

3.3 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. When running an ad-hoc backup, you can back up the data to a vault or to SSI files (safeset image) on disk.

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 an Exchange database, do the following:

- In the **Backup Type** list, do one of the following:
 - To only back up transaction logs and the checkpoint file after the first “seed” backup, click **Incremental**.
 - To back up the database files, checkpoint file and transaction logs, click **Full**.

For more information, see [Plan Full and Incremental Exchange backups](#).

- To validate Exchange data during the backup, select **Validate Exchange database**.

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: Incremental backups for Exchange cannot be deferred, even if deferring is enabled. Deferring can be applied to full backups for Exchange.

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**.

3.4 Plan Full and Incremental Exchange backups

When you run or schedule an Exchange database backup job, you can specify whether to run a Full or Incremental backup. In a Full backup, the database files, checkpoint file and transaction logs are backed up. In an Incremental backup, only the transaction logs and checkpoint file are backed up after the first “seed” backup.

An Incremental backup takes less time to run than a Full backup. However, the time required to recover an Exchange database increases with the number of consecutive Incremental backups. To reduce the amount of time required for a recovery, we recommend performing a Full backup periodically. For example, you could schedule an Exchange backup job to run frequently with the Incremental backup type and periodically (e.g., once per week) with the Full backup type. As shown in the following table, the appropriate backup schedule can also depend on the Exchange database size and traffic.

Exchange database description	Sample backup schedule
Low traffic – approximately 250 users, 4 GB of data, 250 MB of daily data traffic	Full backup every second Saturday night; Incremental backup on other nights
Medium traffic – approximately 1000 users, 16 GB of data, 1 GB of daily data traffic	Full backup every Saturday night; Incremental backup on other nights
High traffic – approximately 4000 users, 64 GB of data, 4 GB of daily data traffic	Full backup every Wednesday and Saturday night; Incremental backup on other nights
High traffic – approximately 4000 users, 64 GB of data, 4 GB of daily data traffic, insufficient bandwidth for large backups during the week	Full backup every Saturday night (which could be deferred to Sunday, if required); Incremental backup on other nights

You should always perform a Full backup after database repair, defragmentation or recovery. These processes significantly change Exchange databases.

Exchange maintenance can affect how much data is transferred during a Full backup. If you run daily maintenance on your Exchange server, the database will change considerably each day. When performing a Full backup, these changes are incorporated into the safeset and will result in longer Full backup times.

When scheduling backup jobs, consider the maintenance window. Backup jobs have priority over mailbox database maintenance. If a backup job runs at the same time as maintenance processes, maintenance will be put on hold until the backup is finished. A maintenance window usually provides enough time for maintenance processes to finish after an Incremental backup, but a Full backup could prevent maintenance processes from running.

3.5 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 Exchange databases

You can restore a Microsoft Exchange database to its original location or to an alternate Exchange database (e.g., a recovery database in Exchange 2010 and later versions). To overwrite an existing database, the database must be unmounted and marked for overwrite.



When restoring an Exchange database, you can specify whether or not to replay transaction logs into the database and mount the database in Exchange. If this option is selected, the logs are rolled forward if they are in the original directory and no log files are missing or corrupt. If this option is not selected, transaction logs are restored to the system but are not replayed into the restored database. The Administrator must review the restored files and manually mount the database.

The process of restoring the database files to the system is recorded in the job logs. The process of replaying transaction logs into the database is recorded in the Windows Event Viewer.

For more information about Exchange restore strategies, see [Restore databases in Exchange 2010 and later versions](#), [Restore Exchange 2007 databases](#) and documentation from Microsoft.


To restore Exchange databases to flat files, see [Restore Exchange databases to flat files](#).

To restore an Exchange database:

1. If you are restoring the database to its original location, delete all files (e.g., .edb, .log and .chk) in the folder where the original database files are located.
Note Do not delete the folder named with a GUID in the name (e.g., 523E7980-EA56-440-9ACB-AF0AE2CF1F0212...).
2. On the navigation bar, click **Computers**.
A grid lists available computers.
3. Find the computer with the Exchange database you want to restore, and expand its view by clicking the computer row.
4. Click the **Jobs** tab.
5. Find the job whose Exchange database you want to restore, and click **Restore** in the **Select Action** menu for the job.
6. In the **Choose What You Want to Restore** dialog box, select **Exchange Databases**, and click **Okay**.
The **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 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: If SSI files were created by a backup to a directory on disk, you cannot restore from the SSI files until they have been imported into the vault and you have synchronized the Agent with the vault.

8. In the **Items to Restore** box, select the check box for each database that you want to restore.
9. Select a **Restore Destination** option:
 - To restore data to the location where it was backed up, select **Restore files to their original location**.
 - To restore to an alternate Exchange database, select **Restore to an alternate Exchange database**. In the **Current Chosen Destination** box, click **Browse**. In the **Select Folder** dialog box, select the location where you want to restore, and click **Okay**.
10. In the **Encryption Password** box, enter the data encryption password. To view the password hint, click the **Hint** button. 
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.
 - Select or clear the **Start Hard Recovery** option.

If the **Start Hard Recovery** option is selected, transaction logs are replayed after the database is restored and the restored database is mounted by Exchange. Logs are also rolled forward if they are in the original directory and no log files are missing or corrupt.

If the **Start Hard Recovery** option is not selected, transaction logs are restored to the system but are not replayed after the database is restored. The restored database is not mounted by Exchange. The Administrator must review the restored Exchange files and manually mount the database.

Click **Okay**.

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.

4.1 Restore databases in Exchange 2010 and later versions

In Exchange Server 2010 and later Exchange versions, you can only restore databases to the active copy in a Database Availability Group. If you restore to the replica copy, you will not be able to mount the database or make it active. You will need to copy the restored files to the active copy node in order to successfully

mount (and take precedence over the other copies). You will also need to update each copy through the Exchange Management Console. For more information, see documentation from Microsoft.

If you are restoring to a new location and you want to mount the database, you must first create a recovery database through the Exchange Management Shell. For more information, see documentation from Microsoft.

To restore an Exchange 2010 or later version database to an alternate location on a non-Exchange server, you must clear the **Start Hard Recovery** option.

If a database's transaction log files are missing or damaged, an incremental backup after the recovery will not succeed. Perform a full backup before you attempt another incremental backup.

4.2 Restore Exchange 2007 databases

In Exchange Server 2007, you can only restore databases to the active node of an Exchange CCR cluster. Otherwise, the restore will fail.

If you are restoring to a new location and you want to mount the database, you must first create a database in a recovery storage group through the Exchange Management Shell. For more information, see documentation from Microsoft.

Restoring a single database to a storage group with more than one database could result in data loss from other databases.

If a database's transaction logs are missing or damaged, an incremental backup after the recovery will not succeed. Perform a full backup before you attempt another incremental backup.

4.3 Restore Exchange databases to flat files


On a computer where the Exchange Plug-in is installed, you can restore an Exchange database to flat files. The Eseutil utility can then be used to bring the data into a database.

To restore an Exchange database to flat files:

1. On the navigation bar, click **Computers**.
A grid lists available computers.
2. Find the computer with the Exchange database 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 Exchange database that you want to restore, and click **Restore** in the **Select Action** menu for the job.
5. In the **Choose What You Want to Restore** dialog box, select **Restore to folder**, and click **Okay**.

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


6. 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 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: If SSI files were created by a backup to a directory on disk, you cannot restore from the SSI files until they have been imported into the vault and you have synchronized the Agent with the vault.

7. In the **Files to Restore** box, select the check box for each database or storage group that you want to restore.

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

9. 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**.

10. To change the log detail level, bandwidth options 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.
- Select or clear the **Start Hard Recovery** option.

If the **Start Hard Recovery** option is selected, transaction logs are replayed after backup data is restored. The restored storage database is prepared for use by Exchange, and logs are rolled forward if they are in the original directory and no log files are missing or corrupt. These processes are recorded in the Windows Event viewer.

If the **Start Hard Recovery** option is not selected, the storage group/database will not be available to Exchange after a restore. The Administrator must review the restored Exchange files and database and manually mount the database. For more information, see documentation from Microsoft.

Click **Okay**.

11. 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.

5 Restore Exchange mailboxes, messages and other objects

If a Microsoft Exchange database is backed up using the Exchange Plug-in, you can restore mailboxes, messages, and other objects from the backup.

To restore items from a Microsoft Exchange backup, you must first use Portal to expose the Exchange safeset as a shared resource. You can then use the Granular Restore for Microsoft Exchange and SQL application to find and restore mailboxes, messages and other Exchange objects.

For more information, or to obtain the Granular Restore for Microsoft Exchange and SQL application, contact your service provider.


To restore Exchange mailboxes, messages and other objects:

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


A grid lists available computers.

2. Find the computer with Exchange objects that you want to restore, and expand its view by clicking the computer row.
3. Click the **Jobs** tab.
4. Find the job with Exchange objects (e.g., messages) that you want to restore, and click **Restore** in the **Select Action** menu for the job.

The **Choose What You Want to Restore** dialog box appears.

5. Select **Mailboxes, messages and other Exchange objects**, and click **Okay**.
6. In the **Restore** dialog box, choose a safeset from which to restore.
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.
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 to the safeset share is now available for you to paste into the Granular Restore application.
12. Launch the Granular Restore for Microsoft Exchange and SQL application. Paste the path to the Exchange safeset share into the Granular Restore application, and then select and restore your data. For more information, see the *Granular Restore for Microsoft Exchange and SQL User Guide*.

13. When you no longer need the safeset share, click **Stop**.

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

6 Monitor computers, jobs 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](#). You can also access logs for unconfigured computers from this page. See [View an unconfigured computers logs](#).
- 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](#).
- Email notifications. To make it easier to monitor backups, users can receive emails when backups finish or fail. See [Monitor backups using email notifications](#).
- 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 Agents.

The **Availability** column indicates whether each Agent 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.
 -  OK with warnings — 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 Agent 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
AppAware	Image	
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.

-  Never Run — Indicates that the backup job has never run.
-  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.
-  Failed — Indicates that the backup failed and no safeset was created.
-  Cancelled
-  Scheduled for Deletion — Indicates that the job is scheduled to be deleted from Portal and job data is scheduled to be deleted from all vaults on the date shown in the Date column. This backup status is only possible in Portal instances where the data deletion feature is enabled.

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

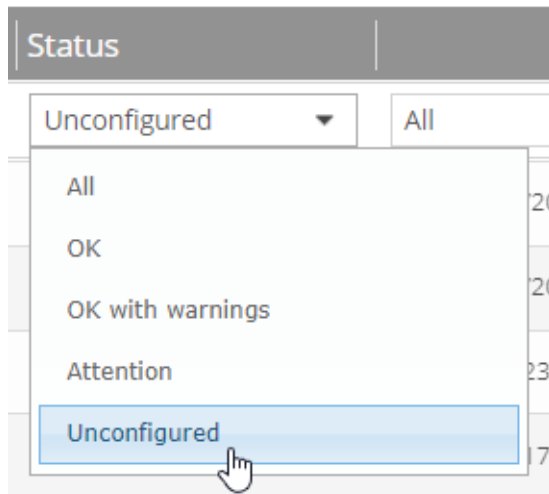
6.2 View an unconfigured computer's logs

You can view logs for unconfigured computers that are online. Unconfigured computers do not have any backup jobs.


To view an unconfigured computer's logs:

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

The Computers page shows registered computers. To only show unconfigured computers, click "Unconfigured" in the **Status** filter.




2. Find an unconfigured computer that is online, and expand its view by clicking the computer row.



3. Click the **logs** link for the unconfigured computer.

The Agent Log Files window shows a list of logs for the computers. Links to the logs appear at the right side of the window.



4. Do one of the following:

- To only view errors and warnings in a log, click **Errors and Warnings** for the log.

- To view an entire log, click **All** for the log.


The log appears in a new browser tab.


```
Log Name: BUAgent-1.XLOG
25-Nov 06:21:49 AGHT-I-04314 Agent Version 8.30.7893 Nov 16 2016 14:12:22
25-Nov 06:21:49 AGHT-I-08103 Executing agent as SYSTEM
25-Nov 06:21:49 AGHT-I-08199 Agent with Id 216bbd19-cbb7-4176-8dfe-be885ee7ecf7 will connect to server qs.corp.com on port 8086
25-Nov 06:21:49 AGHT-I-07466 MIII-4 thread started
25-Nov 06:21:49 AGHT-I-08200 Agent HTTP thread started
25-Nov 06:21:49 AGHT-I-08200 Agent HTTP thread started
25-Nov 06:21:49 AGHT-I-08200 Agent HTTP thread started
25-Nov 06:21:50 AGHT-I-08323 Agent is being redirected to server qs.corp.com on port 8087
25-Nov 06:21:50 AGHT-I-09400 Agent HTTP binding to 127.0.0.1:8031
25-Nov 06:21:50 AGHT-I-09400 Agent HTTP binding to :8031
25-Nov 06:21:54 AGHT-I-07466 MIII-4 thread started
25-Nov 06:21:55 AGHT-E-08307 Failed to set the Agent status to offline.
25-Nov 06:22:01 AGHT-E-08307 Failed to set the Agent status to offline.
25-Nov 06:22:11 AGHT-E-08307 Failed to set the Agent status to offline.
25-Nov 06:22:16 AGHT-I-08914 Agent type set to SERVER
25-Nov 06:22:16 AGHT-E-07514 Failed to Upload System Info in Notification Thread
25-Nov 06:22:21 AGHT-E-07514 Failed to Upload System Info in Notification Thread
25-Nov 06:22:26 AGHT-E-07514 Failed to Upload System Info in Notification Thread
25-Nov 06:22:31 AGHT-E-07477 Failed to Upload Feature Options in Notification Thread
25-Nov 06:22:36 AGHT-E-07477 Failed to Upload Feature Options in Notification Thread
25-Nov 06:22:41 AGHT-E-07477 Failed to Upload Feature Options in Notification Thread
25-Nov 06:22:46 AGHT-E-07476 Failed to Upload Job Types in Notification Thread
```

6.3 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:

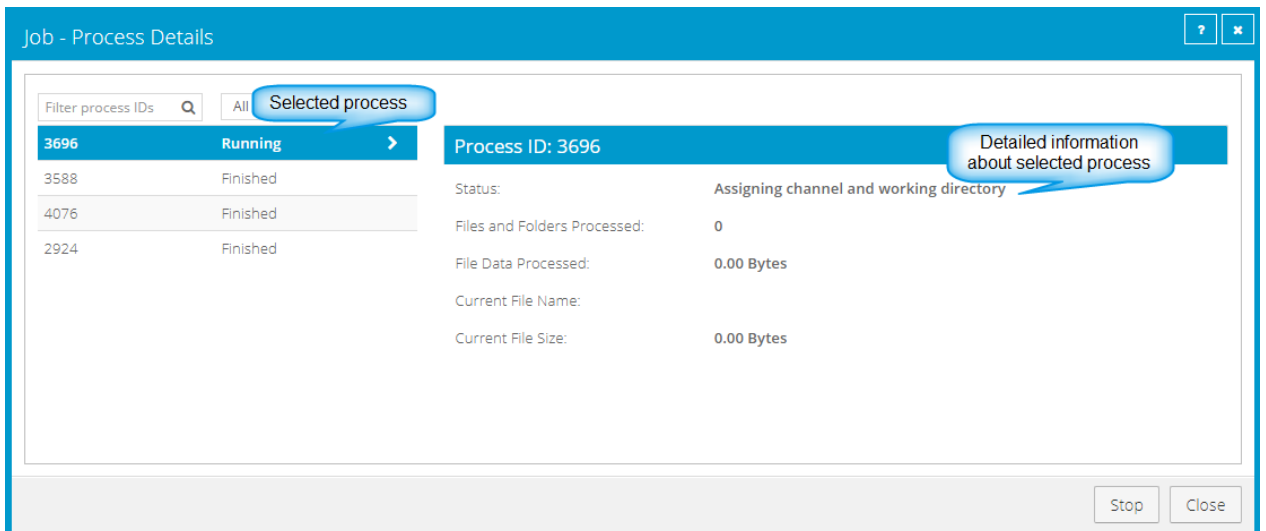
1. Do one of the following:
 - On the Computers page, on the Jobs tab, start a backup, restore or synchronization.
 - On the Computers page, on the Jobs tab, click the “In Progress” symbol  beside the job name.




- 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**.



6.4 Monitor backups using email notifications

To make it easier to monitor backups, users can receive emails when backups finish or fail. Admin users and regular users in Portal can set up email notifications for a computer. See [Set up email notifications for backups on a computer](#).

In some Portal instances, email notifications are configured centrally for Windows systems with Agent version 8.0 or later, instead of separately for each computer. See [Set up email notifications for backups on multiple computers](#).

6.4.1 Set up email notifications for backups on a computer

To set up email notifications for a computer:

1. On the navigation bar, click **Computers**.
2. Find the Agent for which you want to configure email notifications, and click the row to expand its view.
3. On the **Advanced** tab, click the **Notifications** tab.

If the **Notifications** tab appears, but a policy is assigned to the Agent, you cannot change values on the **Notifications** tab. Instead, notifications can only be modified in the policy.

Select one or more of the following checkboxes:

- **On failure.** If selected, users receive an email notification when a backup or restore fails. If a backup fails, you cannot recover any files from the backup.

- **On error.** If selected, users receive an email notification when a backup or restore completes with errors in the log file. You cannot recover files that are backed up with errors, but you can restore other files from the backup (safeset).
- **On successful completion.** If selected, users receive an email notification when a backup or restore completes successfully. You can recover files from a backup that completes, even if there are warnings in the log file.

Email notifications are sent separately for each backup and restore. For example, if three backup jobs fail on a computer and **On failure** is selected for the computer, three notification emails are sent.

If users will receive email notifications after backups and restores, specify the following email notification information:

Email "From" Address	Email address from which email notifications will be sent.
Outgoing Mail Server (SMTP)	Network address of the SMTP that will send the email.
Recipient Address(es)	Email notification recipient email addresses, separated by commas. These should be real, valid email addresses. If one or more is not valid, the transmission to those addresses will fail, and errors will appear in the log files.
Outgoing Server Port (SMTP)	Port number for sending email notifications.
SMTP Credentials	If required, SMTP username, domain, and password.

4. Click **Save**.


6.4.2 Set up email notifications for backups on multiple computers

By default in some Portal instances, Admin users receive emails when backups fail, or are cancelled, deferred, missed or completed. Admin users can select backup statuses for which they want to receive email notifications. These email notifications are sent for Windows systems with Agent version 8.0 or later, instead of separately for each computer.

For other computers, and in Portal instances where Admin users do not automatically receive email notifications, notifications must be configured separately for each computer. See [Set up email notifications for backups on a computer](#).

To set up email notifications for backups on multiple computers:

1. When signed in as an Admin user, click your email address at the top right of the Portal page.
The user menu appears.



2. Click **Profile Settings**.

Your user profile appears. If your profile includes an Email Notification Settings section with a list of backup events (e.g., Backup Canceled, Backup Completed), you can select events for which you want to receive emails.

Email Notification Settings

We can keep you updated via email on the status of backup jobs for your site. Select each backup status for which you want to receive email notifications.

- Backup Cancelled
- Backup Completed
- Backup Completed with Errors
- Backup Completed with Warnings
- Backup Deferred
- Backup Failed
- Backup Missed
- Encryption Password Changed

Update notifications

If Email Notifications Settings do not appear, you must set up notifications separately for each computer. See [Set up email notifications for backups on a computer](#).

3. In the Email Notification Settings list, select any of the following events for which you want to receive emails:
 - Backup Cancelled
 - Backup Completed
 - Backup Completed with Errors
 - Backup Completed with Warnings
 - Backup Deferred
 - Backup Failed
 - Backup Missed
4. Click **Update notifications**.

6.5 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 Agents.

2. Find the Agent 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.

Name	Job Type	Description	Last Backup Status	Date	Action
BMRJob	Local System		✔ Completed	today at 9:31 AM	Select Action
CloudServerBackup	Local System	This backup protects your entire C drive. It will be backed up to the cloud, per your retention schedule.	✔ Completed	yesterday at 7:32 PM	Select Action
Job	Local System		⚠ Completed with warnings	yesterday at 10:45 PM	Select Action

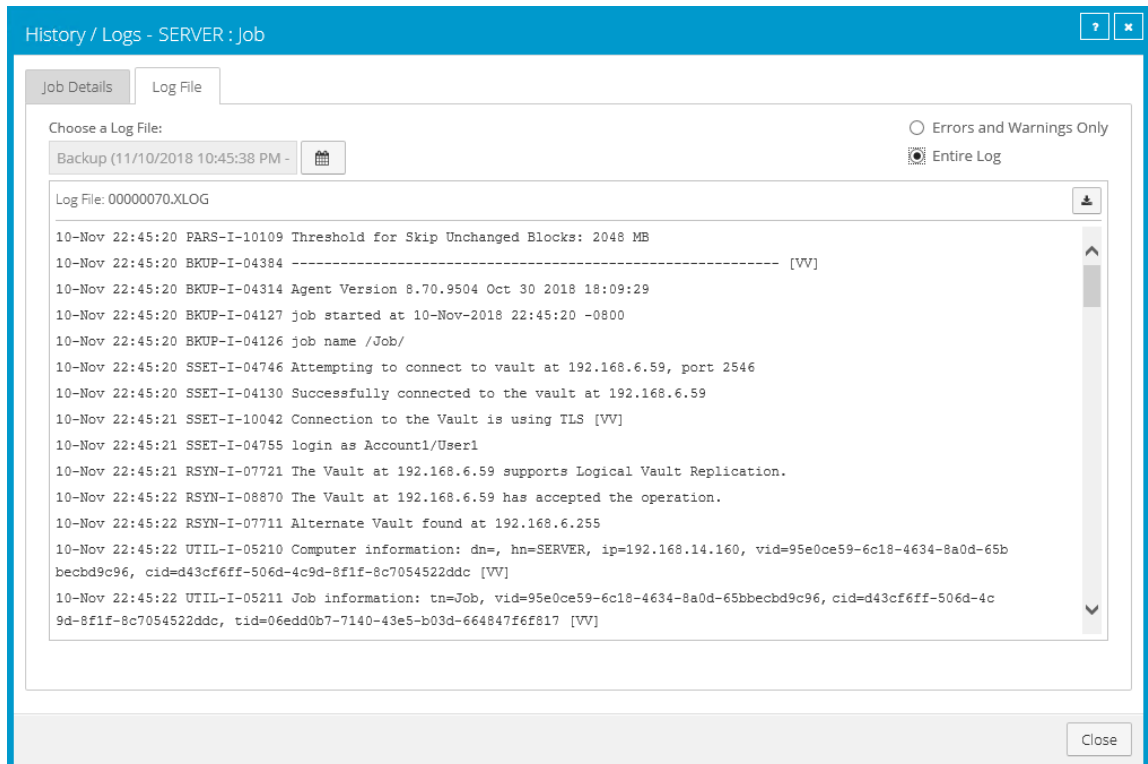
3. To view log files for a job, do one of the following:

- 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.

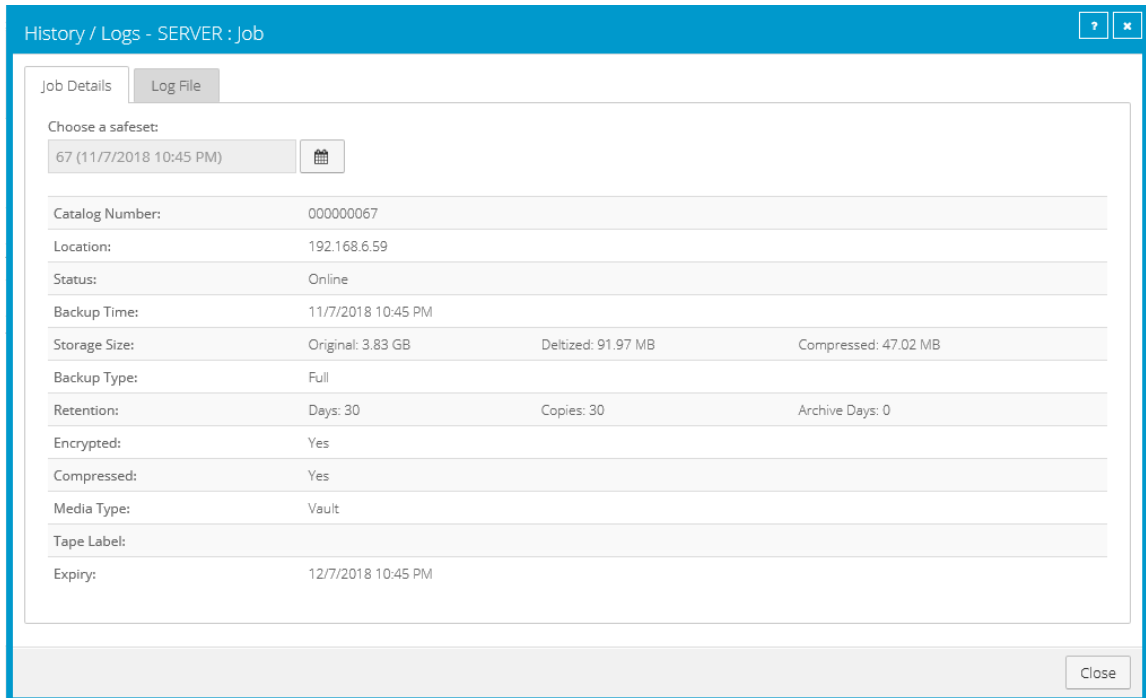
4. To view processes for a different day, click the calendar button. In the calendar that appears, click the date of the log that you want to view.

- 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.



- 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 icon. 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.



6.6 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.

The screenshot shows the "Monitor" page with a "Display: Jobs" dropdown. At the top right, there are controls for "Export" (dropdown), "Show 25 Records per page", "Save View", and "All Jobs" (dropdown). Below these is a table with the following columns: Computer, Job Name, Last Backup Status, Date, Backup Size, and Site Name.

Computer	Job Name	Last Backup Status	Date	Backup Size	Site Name
		All			
WINDOWSSERVER	restjob	Completed with warnings	yesterday at 7:30 PM	281.99 MB	Site1
WINDOWSSERVER	Backup1	Completed with warnings	yesterday at 9:00 PM	7.40 MB	Site1
PROTECTEDSERVER	BMRjob	Completed	today at 9:31 AM	21.21 GB	Site1
WINDOWSSERVER	CloudServerBackup	Failed	yesterday at 10:45 PM	0.00 Bytes	Site1
PROTECTEDSERVER	CloudServerBackup	Completed	yesterday at 7:32 PM	17.91 GB	Site1
PROTECTEDSERVER	Job	Completed with warnings	yesterday at 10:45 PM	3.84 GB	Site1

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.