

Selecting location

The **Locations** tab of the Backup Policy wizard allows you to specify the backup location for each operation you include in the backup policy.

For each operation you have included in the backup policy, you can specify the location type, full path in which to store the backup file, an optional housecleaning schedule for existing disk archives, and the backup file extension.

What types of backup locations can you use?

SQL Safe supports the following location types:

- Back up to a single file on the local computer or a network share
- Back up to multiple striped files on the local computer or a network share
- Back up to tape using Tivoli Storage Manager
- Back up to Data Domain
- Back up to Amazon S3 Cloud
- Back up to Microsoft Azure Cloud
- Back up to tape using Tivoli Storage Manager Striped Files

What do you do if you do not have an existing archive?

If you do not specify an existing archive, SQL Safe creates a new backup set with the name you specify. The location entered for each backup type must be valid for all SQL Server instances. You can choose to **Append** or **Overwrite** if the archive already exists.

What accounts can you specify to access the backup files location?

Specify the account to access the filesystem when you select any of the following location types: Single File, Striped files, or Data Domain. Depending whether you selected the SQL Server Agent or the SQL Safe Backup Agent for your policy, you will be able to choose between SQL Server Agent service account/SQL Safe Backup Agent service account respectively or a Windows account. Click **Account** and select your preferred option.



The account specified must have read and write privileges on the directory selected for your backup file location.

How do you keep the backups running despite network errors?

Select **Enable network resiliency** and then click **Configure** to change the default settings. By default, SQL Safe will retry the backup operation every 10 seconds and fail after 5 minutes (300 seconds) of continuous errors. Over the course of the operation, SQL Safe allows a total of 60 minutes in which to retry the backup before stopping the operation.

This option is not available when backing up to tape using Tivoli Storage Manager or Amazon S3 Cloud.

Can you change the default file locations?

SQL Safe automatically populates the path using several available variables, depending on the location type. You can modify this path to suit your needs, taking advantage of all the available variables.

For a disk backup, browse for or type the directory in which to store the backup file. You can use the supplied macros in the way best suited to your storage needs. And if you want to limit the lifetime of your backup sets created by the policy, you can select the option that removes files older than the specified time.

For a TSM backup, browse for or enter the high level directory for the tape file. Then, browse for or enter the location of the TSM configuration file.

What does removing old files do?

For backups written to a single file or mirrored files, you can choose to remove old files to prevent disk space limitations. When you select to remove files older than the specified time, backup files created with names of the same format will be deleted from that directory. You can configure SQL Safe to delete old backup files from the primary archive as well as from your mirror archives.

For backups written to a TSM Server, you can configure SQL Safe to mark these files as inactive after a specified age. This option is not available when using Amazon S3 Cloud.

How do you mirror the backups this policy creates?

When selecting Single File or Data Domain as location types, you have the option to create mirror archives. For each mirror archive, SQL Safe creates a copy of the backup archive set. Click **Mirror Archives** and specify where you want the mirrored files to be stored. Take into account that you can specify up to two mirrors for each backup operation.

If you want to stop the backup operation when mirror location is unavailable, select **Abort backup if a mirror location reports a failure**. You can also enable to remove files older than a specified time.



Keep in mind that creating mirrors can impact the performance of your backup operation.

What do you specify when backing up to a TSM Server?

When a TSM location is selected, you have to specify the following settings:

- The location of the TSM Client configuration file
- High level file path
- Low level file path
- Management Class

You can also configure SQL Safe to mark these files as inactive after a specified age.



SQL virtual database is not available when backing up to a TSM Server.

How do you backup to multiple stripes using TSM Servers?

You can also back up to tape using Tivoli Storage Manager Striped Files. If you select this option, provide the following information:

- The location of the TSM Client configuration file
- High level file path
- Management class
- The number of striped archives to use and the respective names for each low level file name

You can also configure SQL Safe to mark these files as inactive after a specified age.



Take into account:

- If the number of stripes is greater than the available sessions on TSM server, the backup fails with a message "sessions are not available on TSM". There is no available way for the TSM client to find out available sessions on the TSM server.
- If a password is required for TSM, then policies created using TSM may not work correctly.

What do you specify when backing up to Amazon S3 Cloud?

When Amazon S3 Cloud backup location is selected, you have to specify the following fields:

- **Append/Overwrite** - select if you want to append the backup archive to an existing one or if you prefer to overwrite it.
- **Bucket Name** - define the name of the Amazon S3 bucket where your backup will be stored.
- **Secret Key** - specify the secret key generated in your security credentials of your Amazon S3 web console. For more information, click [here](#)
- **Access Key** - specify the access key generated in your security credentials of your Amazon S3 web console. For more information, click [here](#).
- **File size** - File size is used to determine the minimal parts of the backup file in bytes that will be sent to the bucket simultaneously. The minimal value for File Size is 1 000 000 bytes. Note that when backing up to a cloud location, the network quality may affect performance.
- **Region** - select the region where your information will be stored. You can find more information about these regions [here](#).
- **Filename** - specify the name with which your information will be stored.

What do you specify when backing up to Microsoft Azure Storage?

When Microsoft Azure Storage is selected, make sure you specify the following fields:

- **Container name** - the name of the Azure container where your backup will be stored. Every [blob](#) in Microsoft Azure storage must reside in a container. The container forms part of the blob name. For more information, click [here](#).
- **Azure Storage Account Name** - the account name of your storage account. Every object you store in Azure Storage has a unique URL address. The storage account name forms the subdomain of that address. You can find more information in the following [link](#).
- **Azure Access Key** - the access key to your Azure Storage Account. For more information about Azure Keys, click [here](#).
- **Filename** - to be used as your primary backup archive.

After defining your setting in this tab, click **Next** to [Configure the schedule of your backup operations](#).

SQL [Safe](#) is a high-performance backup and recovery solution for your SQL Servers. [Learn more](#) > >

IDERA Website	Products	Purchase	Support	Community	About Us	Resources	Legal
-------------------------------	--------------------------	--------------------------	-------------------------	---------------------------	--------------------------	---------------------------	-----------------------