

# USING SOFT DELETE FOR BLOB STORAGE

One of the most useful and valuable features inside Azure's blob storage portfolio is the ability to enable soft delete, which keeps a copy of the deleted data for a specified period of time, mitigating against 90% of the cases where an accidental deletion of data occurs.

For optimal protection for your blob data, Microsoft recommends enabling:

- Container soft delete, to restore a container that has been deleted.
- Blob versioning, to automatically maintain previous versions of a blob.
- Blob soft delete, to restore a blob, snapshot, or version that has been deleted.

## HOW DOES IT WORK?

When you enable blob soft delete for a storage account, you specify a retention period for deleted objects of between 1 and 365 days. The retention period indicates how long the data remains available after it's deleted or overwritten. The clock starts on the retention period as soon as an object is deleted or overwritten.

While the retention period is active, you can restore a deleted blob, together with its snapshots, or a deleted version.

You can change the soft delete retention period at any time. An updated retention period applies only to data that was deleted after the retention period was changed. Any data that was deleted before the retention period was changed is subject to the retention period that was in effect when it was deleted.

Attempting to delete a soft-deleted object doesn't affect its expiry time.

If a blob has snapshots, the blob can't be deleted unless the snapshots are also deleted. When you delete a blob and its snapshots, both the blob and snapshots are marked as soft-deleted. And no new snapshots are created.

You can also delete one or more active snapshots without deleting the base blob. In this case, the snapshot is soft deleted.

If a directory is deleted in an account that has the hierarchical namespace feature enabled on it, the directory and all its contents are marked as soft deleted. And only the soft deleted directory can be accessed. In order to access the contents of the soft-deleted directory, the soft-deleted directory needs to be undeleted first.

Soft-deleted objects are invisible unless they're explicitly displayed or listed.

## CATERING FOR OVERWRITES OF DATA

When blob soft delete is enabled, overwriting a blob automatically creates a soft-deleted snapshot of the blob's state prior to the write operation. When the retention period expires, the soft-deleted snapshot is permanently deleted.

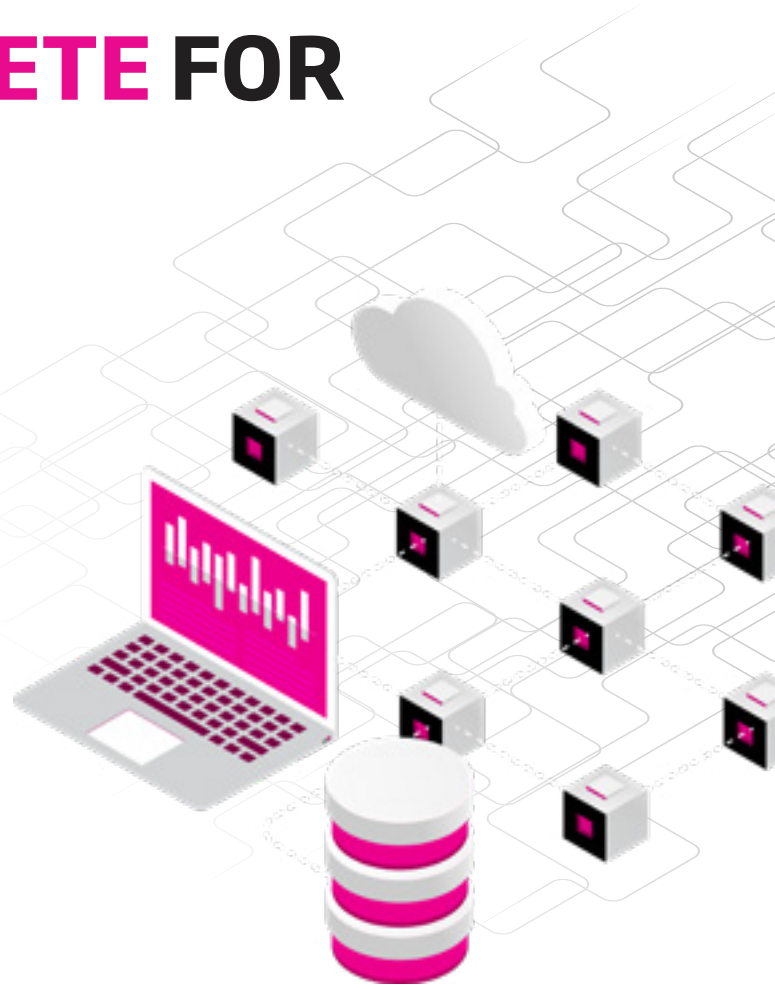
Soft-deleted snapshots are invisible unless soft-deleted objects are explicitly displayed or listed.

To protect a copy operation, blob soft delete must be enabled for the destination storage account.

Blob soft delete doesn't protect against operations to write blob metadata or properties, however. No soft-deleted snapshot is created when a blob's metadata or properties are updated.

Blob soft delete doesn't afford overwrite protection for blobs in the archive tier. If a blob in the archive tier is overwritten with a new blob in any tier, then the overwritten blob is permanently deleted.

And what's great is, for premium storage accounts, soft-deleted snapshots don't count toward the per-blob limit of 100 snapshots.



## RESTORING SOFT DELETED OBJECTS

You can restore soft-deleted blobs or directories (in a hierarchical namespace) within the retention period.

The Undelete Blob operation restores a blob and any soft-deleted snapshots associated with it. Any snapshots that were deleted during the retention period are restored. In accounts that have a hierarchical namespace, the access control list of a blob is restored along with the blob.

In accounts that have a hierarchical namespace, the Undelete Blob operation can also be used to restore a soft-deleted directory and all its contents.

That said, if you rename a directory that contains soft-deleted blobs, those soft-deleted blobs become disconnected from the directory. If you want to restore those blobs, you'll have to revert the name of the directory back to its original name or create a separate directory that uses the original directory name. Otherwise, you'll receive an error when you attempt to restore those soft-deleted blobs.

You also cannot restore a directory or a blob to a filepath that has a directory or blob of that name already there.

You cannot access the contents of a soft-deleted directory until after the directory has been undeleted.

## VERSIONING

If blob versioning and blob soft delete are both enabled for a storage account, then overwriting a blob automatically creates a new previous version that reflects the blob's state before the write operation.

The new version isn't soft-deleted and isn't removed when the soft-delete retention period expires. No soft-deleted snapshots are created.

If blob versioning and blob soft delete are both enabled for a storage account, then when you delete a blob, the current version of the blob becomes a previous version, and there's no longer a current version. No new version is created and no soft-deleted snapshots are created. All previous versions are retained until they are explicitly deleted, either with a direct delete operation or via a lifecycle management policy.

Enabling soft delete and versioning together protects previous blob versions as well as current versions from deletion. When soft delete is enabled, explicitly deleting a previous version creates a soft-deleted version that is retained until the soft-delete retention period elapses. After the soft-delete retention period has elapsed, the soft-deleted blob version is permanently deleted.

You can use the Undelete Blob operation to restore soft-deleted versions during the soft-delete retention period. The Undelete Blob operation always restores all soft-deleted versions of the blob. It isn't possible to restore only a single soft-deleted version.



## CONSIDER COSTS

All soft-deleted data is billed at the same rate as active data. And you won't be charged for data that is permanently deleted after the retention period elapses.

When you enable soft delete, it's recommended that you choose a short retention period to better understand how the feature will affect your bill. The minimum recommended retention period is seven days.

Enabling soft delete for frequently overwritten data may result in increased storage capacity charges and increased latency when listing blobs. You can mitigate this additional cost and latency by storing the frequently overwritten data in a separate storage account where soft delete is disabled.

You aren't billed for transactions related to the automatic generation of snapshots or versions when a blob is overwritten or deleted. You're billed for calls to the Undelete Blob operation at the transaction rate for write operations.