

Feature

Trash translator for GlusterFS

Summary

This feature will enable user to temporarily store deleted files from GlusterFS for a specified time period.

Current Status

In the present scenario deletion by a user results in permanent removal of a file from the storage pool. An incompatible translator code for trash is currently available as part of codebase. On the other side gluster cli lacks a volume set option to load the trash translator in volume graph.

Detailed Description

Trash is a desired feature for users who accidentally delete some files and may need to get back those in near future. Currently, GlusterFS codebase includes a translator for trash which is not compatible with the current version and so is not usable by users. Trash feature is planned to be implemented as a separate directory in every single brick inside a volume. This would be achieved by a volume set option from gluster cli.

A file can only be deleted when all hard links to it has been completely removed. This feature can be extended to operations like truncation where we need to retain the original file.

Benefit to GlusterFS

With the implementation of trash, accidental deletion of files can be easily avoided.

Scope

- *Nature of proposed change*

Proposed implementation mostly involves modifications to existing code for trash translator.

- *Implications on manageability*

Gluster cli will provide an option for creating trash directories on various bricks.

- *Implications on presentation layer*

None

- *Implications on persistence layer*

None

- *Implications on 'GlusterFS' backend*

The overall brick structure will include a separate section for trash in which regular files will not be stored, i.e. space occupied by the trash become unusable.

- *Modification to 'GlusterFS' metadata*

The original path of files can be stored as an extended attribute.

- *Implications on 'glusterd'*

An alert can be triggered when trash exceeds a particular size limit. Purging of a file from trash depends on its size and age attributes or other policies.

- *Implications on Rebalancing*

Trash can act as an intermediate storage when a file is moved from one brick to another during rebalancing of volumes.

- *Implications on Self-healing*

Self-healing must avoid the chance of re-creating a file which was deleted from a brick while one among the other bricks were offline. Trash can be used to track the deleted file inside a brick.

- *Scope of Recovery*

This feature can enhance the restoring of files to previous locations through gluster cli with the help of extended attributes residing along with the file.

How to Test

Functionality of this trash translator can be checked using self-heal and rebalancing options.

User Experience

Users can access files which were deleted accidentally or intentionally and can review the original file which was truncated.

Dependencies

None

Documentation

Status

In design

Comments and Discussion
