

IBM Spectrum Scale (GPFS) Snapshots

The cluster uses Spectrum Scale (GPFS) to provide highspeed shared filesystem access between cluster nodes. Spectrum Scale includes the ability to create filesystem snapshots which create temporary backups of data stored in the filesystem.

This document outlines how users can use these snapshots to recover their data in case of accidental deletion or other changes.

Snapshot Schedule

In the cluster environment the Spectrum Scale software is setup to record snapshots of the /gpfs/home and /gpfs/projects folders using the following schedule.

Location	Days to keep	Months to keep
/gpfs/home	7	3
/gpfs/projects	7	3

The snapshot schedule maintains 7 days of snapshots and one snapshot from the last 3 months. Snapshots are deleted automatically after they move outside this schedule.

Recovering Data from Snapshots

A hidden folder called `.snapshots` exists in the `/gpfs/home` and `/gpfs/projects` folders. If you access the `.snapshots` folder you will see a directory listing similar to the output below.

```
[user@hn1 .snapshots]$ pwd
/gpfs/home/.snapshots
[user@hn1 .snapshots]$ ls -la
total 6
dr-xr-xr-x 2 root root 8192 Dec 31 1969 .
drwxr-xr-x 20 root root 4096 Jun 29 14:20 ..
drwxr-xr-x 15 root root 4096 Apr 28 13:33 @GMT-2022.06.01-03.00.37
drwxr-xr-x 20 root root 4096 Jun 29 14:20 @GMT-2022.07.01-03.00.56
drwxr-xr-x 20 root root 4096 Jun 29 14:20 @GMT-2022.07.28-03.00.25
drwxr-xr-x 20 root root 4096 Jun 29 14:20 @GMT-2022.07.29-03.00.26
drwxr-xr-x 20 root root 4096 Jun 29 14:20 @GMT-2022.07.30-03.00.28
drwxr-xr-x 20 root root 4096 Jun 29 14:20 @GMT-2022.07.31-03.00.30
drwxr-xr-x 20 root root 4096 Jun 29 14:20 @GMT-2022.08.01-03.00.30
drwxr-xr-x 20 root root 4096 Jun 29 14:20 @GMT-2022.08.02-03.00.31
drwxr-xr-x 20 root root 4096 Jun 29 14:20 @GMT-2022.08.03-03.00.33
```

Each of the dated folders (eg- `/gpfs/home/.snapshots/@GMT-2022.08.03-03.00.33`) contains a copy of the directory structure as it existed on the date the snapshot was taken.

To recover files from the snapshot simply browse the directory structure in the snapshot folder to locate the files then copy them to the desired restore point.

Permissions are preserved in the snapshot so data access restrictions are the same as the original source. All users can manually recover any data that they had access to in the original source location.