PANZURA

SOLUTION BRIEF

Permanent, Secure File Erasure

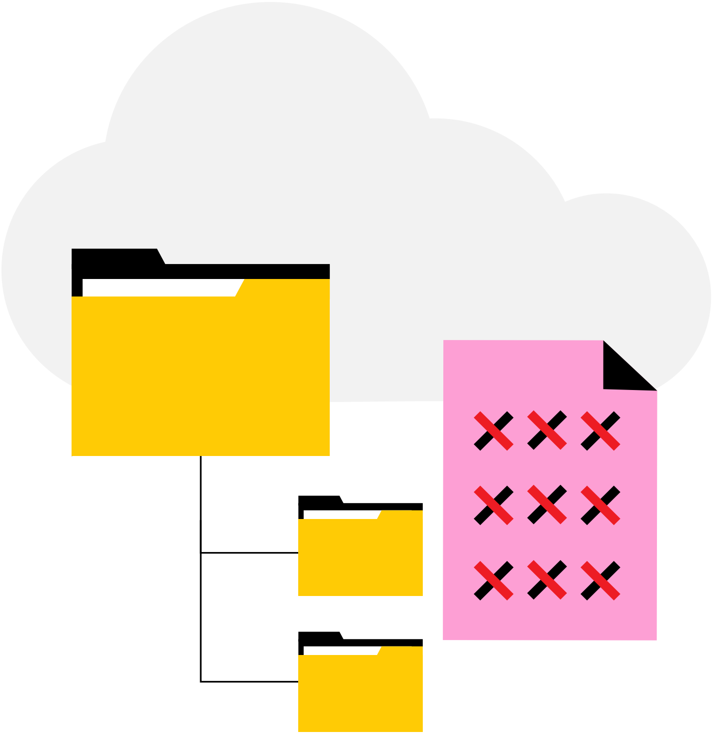

Panzura's global file system CloudFS creates data resiliency against accidental deletion, ransomware attacks or other damage by making it immutable as it's moved to object storage. Data is then further protected by read-only snapshots taken at – and kept for – configurable intervals.

However, CloudFS also supports the ability to completely and securely delete sensitive files when required, permanently erasing them in such a way that they cannot be recovered.

For organizations running legacy storage along with regular backups and offsite archival processes, this is a time-consuming process that requires identifying and deleting every instance of a file in primary storage as well as in backups and offsite storage.

Not only is this problematic for older files, but this approach doesn't guarantee that all traces of the file have been removed, or that it cannot be restored.

Panzura CloudFS Secure Erase makes it possible to permanently erase a file or folder. A successful Secure Erase operation removes all versions of specified files and folders from the file system, and its local nodes as well as the associated objects stored in the cloud.

Understanding How CloudFS Stores Data

To understand how Secure Erase works, let's first take a look at how Panzura CloudFS stores data.

As files are created and edited by users, files are written to the local Panzura CloudFS node they are working from, and CloudFS generates and adds metadata – information about the file.

The file data and metadata are then deduplicated, compressed, encrypted, and split into 128KB blocks. CloudFS deduplicates data on the fly – before moving it into object storage, to avoid storing redundant copies of identical data.

Where identical data is found within the 128KB data block, it is not saved to object storage. Instead, metadata pointers for the file are updated to point to the original block containing that identical data.

At the same time, the data reference count ("refcnt") of that data block is incremented to track the number of times it is referenced within CloudFS.

Unique data blocks are then packaged together into Drive Files of 4MB each and these Drive Files are transported to the object store. The process of transporting new data to the object store and syncing changes throughout CloudFS runs every 60 seconds, on every node in CloudFS.

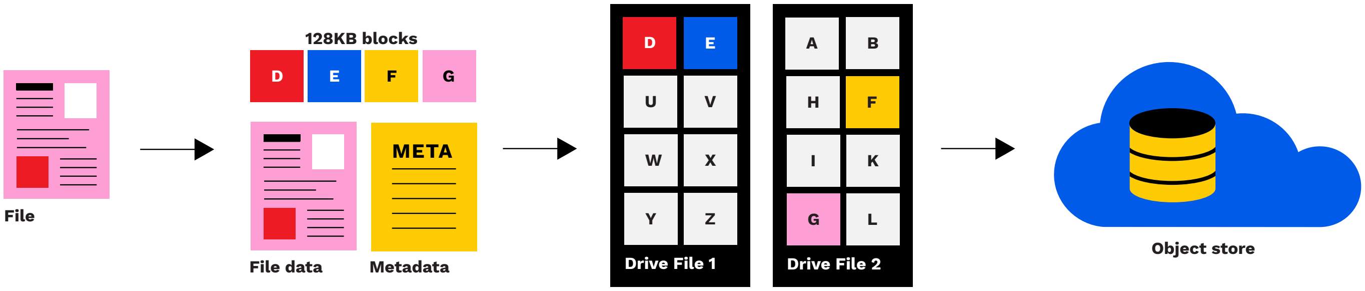

CloudFS Secure Erase

Now that we understand how data is stored, let's look at how the Secure Erase process works on CloudFS.

The Secure Erase process is first initiated on the CloudFS node that owns the file – that is, the node that currently holds the file lock. This is the node on which the file was last opened. With CloudFS, a file must be removed from three different locations in order to be securely erased from your file system.

Erasure from Local CloudFS nodes

When the Secure Erase process is initiated on a file, CloudFS identifies all the blocks (and their respective locations on the local CloudFS node) which form the metadata and data of the file.

The drive files containing these blocks are located and downloaded to the local CloudFS memory.

All metadata blocks associated with the file are deleted from the active file system.

After this step, there is no way of recovering the file from CloudFS because the metadata contains all the information required to retrieve user data blocks.

Next, all unique data blocks which comprise the file are then overwritten with zeros, replacing the original content with meaningless data.

This is a second layer of assurance that data is not recoverable by the file system.

Erasure from the Object Store
The drive files containing these overwritten data blocks are then transported to the object store as part of the normal write and sync process that runs every 60 seconds.

The metadata updates sent to the object store, and to every location in the file system, within that same sync process now ensure that there is no remaining metadata record of the erased file in any location.

Erasure from Snapshots
Users can create a snapshot either manually, or through the CloudFS snapshot scheduler.

Snapshots consist of both a listing of files and folders as well as references to the data blocks which comprise those respective files and folders. In order for a file to be securely erased, references to a file need to be removed and snapshots referring to the file need to be deleted.

When performing Secure Erase on a file, CloudFS will identify and list the manually created snapshots with blocks that reference the target file.

CloudFS will not alter the snapshots in any way – users must take proactive steps to remove these snapshots.

Note: Any scheduled snapshot must age out. These cannot be actively removed from the CloudFS system.

Scheduled snapshots are configurable and are retained according to a schedule set by each organization, to meet their requirements.

For example, a weekly snapshot is taken and four weekly snapshots are kept. The oldest weekly snapshot will age out once four new weekly snapshots are taken.

CloudFS requires the use of Master Snapshots for internal operations. These snapshots cannot be removed by the user and will age out of the system, usually within weeks.

While retained, these snapshots still contain references to the erased file. However, if the snapshot references data that has been securely overwritten, the file will not be accessible or readable in any way.

Confirmation of File Erasure
Confirmation of completion of the secure erase process can be obtained as follows:

Proceed to the Secure Erase console through the CloudFS UI (Home > Maintenance > CloudFS Operations > Secure Erase). Under the “Files” section, click on the filename of the target file. Confirm that the status of the file shows “SECURELY ERASED”.

Furthermore, selecting “Download Report” will show the actions taken and related timestamps, to confirm a completed Secure Erase process.

Example download report:

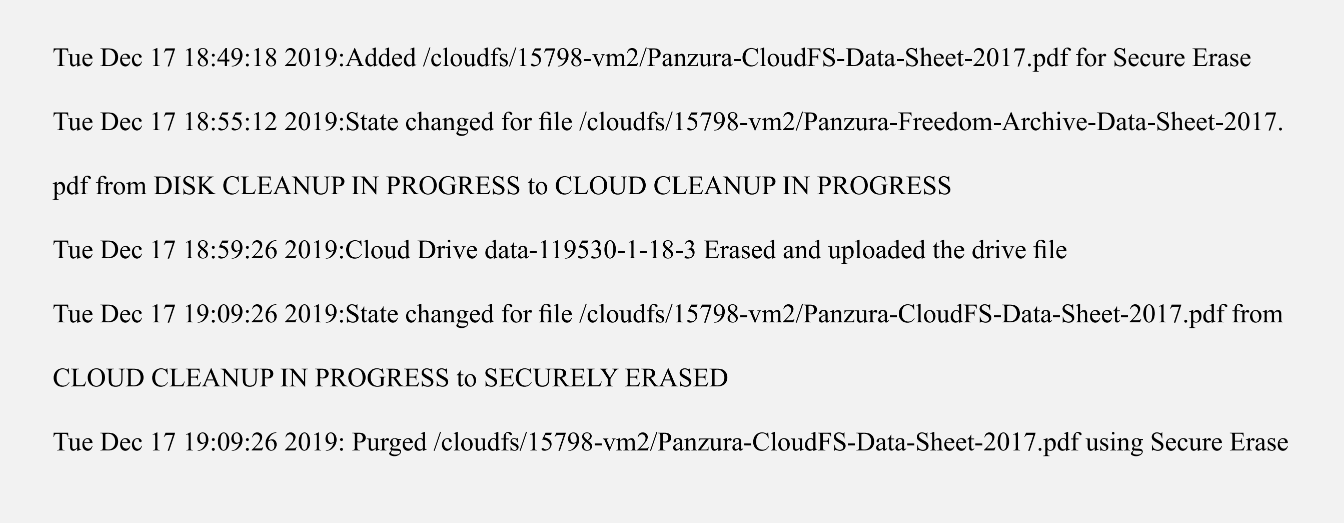

```
Tue Dec 17 18:49:18 2019:Added /cloudfs/15798-vm2/Panzura-CloudFS-Data-Sheet-2017.pdf for Secure Erase

Tue Dec 17 18:55:12 2019:State changed for file /cloudfs/15798-vm2/Panzura-Freedom-Archive-Data-Sheet-2017.

pdf from DISK CLEANUP IN PROGRESS to CLOUD CLEANUP IN PROGRESS

Tue Dec 17 18:59:26 2019:Cloud Drive data-119530-1-18-3 Erased and uploaded the drive file

Tue Dec 17 19:09:26 2019:State changed for file /cloudfs/15798-vm2/Panzura-CloudFS-Data-Sheet-2017.pdf from

CLOUD CLEANUP IN PROGRESS to SECURELY ERASED

Tue Dec 17 19:09:26 2019: Purged /cloudfs/15798-vm2/Panzura-CloudFS-Data-Sheet-2017.pdf using Secure Erase
```

Comprehensive Support for Secure Deletion of Sensitive Files

For IT environments that require the ability to securely remove all traces of highly sensitive files, CloudFS Secure Erase makes it possible to delete a file or folder so that the contents cannot be restored, even using the most advanced technology available.

CloudFS secure erase is the highest purge level that can be attained without physically destroying the disk drives. It removes all versions of specified files and folders from the Panzura node and the associated objects stored in the cloud.

Not only are file references removed, but all data within the object store are securely erased and replaced with zeros, making the original file irretrievable and unreadable.

Secure erase can be used with any supported cloud provider.

About Panzura

Panzura makes hybrid multi-cloud data management seem easy. Panzura's multi-cloud data management platform is a single, unified data engine designed to securely power the most rigorous, large-scale multi-site enterprise data workflows across the globe. Intelligent edge technologies enable LAN performance with cloud economics together with simplified data management, advanced analytics, reduced operational complexity, and improved security.