


# Stay Up, When Your Storage Goes Down

Cloud Mirroring: The Ultimate High Availability Solution for Enterprise Data.





## Stay Up, Even When Your Primary Storage Goes Down.

When data is critical to powering your organization, it's paramount that your files are available whenever they're needed, wherever they're needed, to any authorized user or process looking for them.

Redundancy within your IT stack is as critical as it has ever been, and as experience has shown, even the giant public cloud storage providers can experience outages. Infrequent though they may be, their disruptive nature can prove exceptionally costly.

CloudFS cloud mirroring lets you place the same set of data in two separate object stores, giving you data redundancy with high availability. If your primary cloud goes down, CloudFS switches to your secondary cloud store, with no interruption, and no data loss, letting you carry on while your primary cloud provider is offline.

# Achieving Multi-Cloud Storage and File Operation Redundancy

Enterprise-grade modern cloud storage solutions are typically engineered to offer high levels of data durability, offering features such as geographic replication to make data resilient in the long term. However, data availability metrics are often substantially less robust, while the impact to business processes that rely on immediate availability of data is significant when that data is inaccessible.

**Panzura cloud mirroring delivers high data availability by maintaining an exactly mirrored data set in a secondary object store, providing uninterrupted service in the case of a cloud storage outage.**

The 8th generation of Panzura's award-winning global cloud file system CloudFS™ introduces the ability to mirror data across two separate object stores, writing to both simultaneously in real time. With cloud mirroring enabled, you can operate both a primary and secondary object store, with both holding an identical data set at all times. A real time write-split captures new and changed data in both object stores, as new files are created, or edits are made.

If your primary cloud store suffers a outage, CloudFS will automatically failover to the secondary store, allowing front-end file services for systems or users to continue to operate. When the primary cloud object store is available again and you fail back, Panzura will automatically synchronize both clouds to a consistent data set.

Not only does cloud mirroring eliminate your dependency on a single cloud or object storage vendor, this multi-cloud solution protects you against disruption and data loss from accidental cloud bucket deletion, as well as cyber threats against your cloud provider.

CloudFS ensures that both cloud stores are completely consistent, holding the exact same set of data down to the last byte, at all times. If the event you need to failover to the secondary store, every single file is there, accessible to every authorized user.

# Cloud Mirroring in Operation

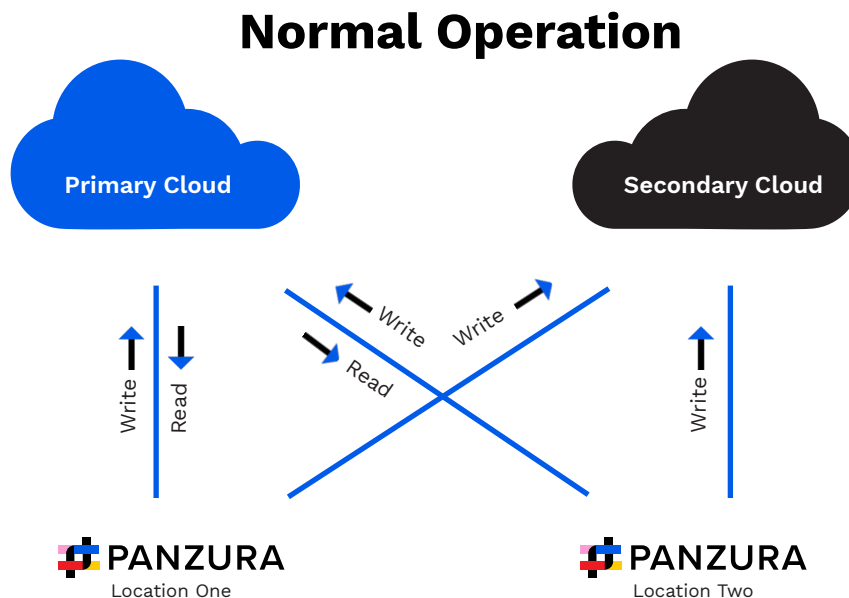
CloudFS uses a cloud connector to communicate with any compatible object store via that cloud’s RESTful API. This object store may be a public cloud, private (on premises) cloud, or completely “dark” cloud with no external connection.

Any two object stores compatible with CloudFS can be designated your primary and secondary clouds. These may be two separate cloud regions from the same provider, two data tiers from the same provider, or two distinct object stores from unrelated providers. Typically, the secondary object store carries a lesser cost than the primary. Each store has its own, independent cloud connector.

Each location in a CloudFS file network reads from the primary cloud in real time, anticipating and caching the most used files locally for high performance.

Every 60 seconds, every location in the CloudFS network simultaneously writes new and changed data to both the primary and secondary clouds, storing it as immutable data. This write-split ensures a complete, redundant copy of any additional data is captured in the secondary cloud, as well as being securely stored in the primary cloud, and available to every other location for immediate data consistency.

This diagram illustrates a two-location CloudFS under normal operation, with cloud mirroring activated. Every location reads from and writes to the primary cloud, and writes to the secondary cloud in parallel.



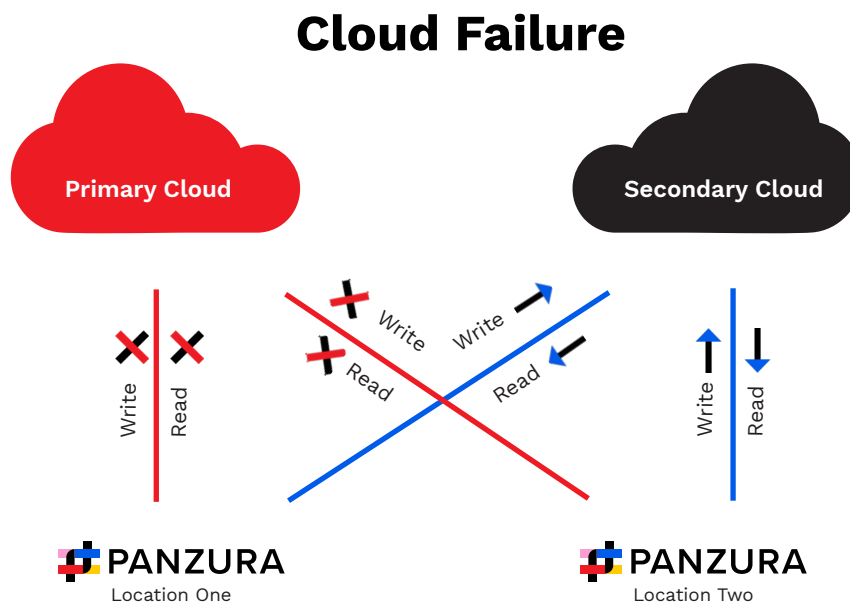
# Cloud Failure and Failover

In the event of a cloud failure, read and write operations to the failed cloud are disabled. With cloud mirroring activated, a sustained primary cloud outage will result in CloudFS failing over to the secondary cloud for read and write operations, until the primary cloud is restored.

Because every location writes to both clouds simultaneously, data held in the secondary cloud is completely consistent with the data in the primary cloud, resulting in a seamless experience with no data loss or file inconsistencies, even in the event of a catastrophic cloud storage failure.

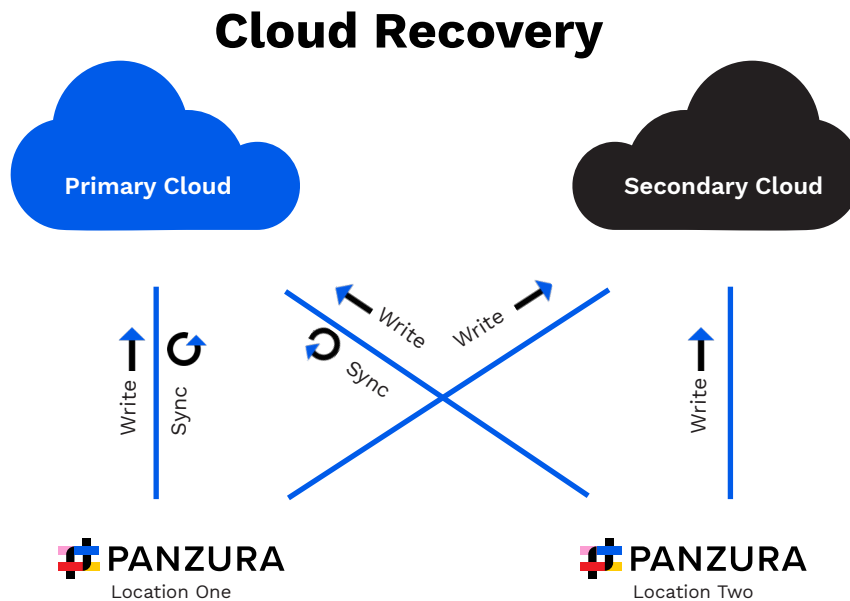
CloudFS does not depend on the cloud for any file operations. As a result, a cloud failure does not impact on functions such as file locking, or immediate data consistency between locations.

As shown below, while the primary cloud is unavailable, every location in a CloudFS reads from and writes to the secondary cloud. This allows file operations to continue uninterrupted, while maintaining a record of data created locally that has not been synced to the primary cloud.



When the primary cloud is available again, CloudFS will resume writing to both clouds in parallel, and an administrator can switch read operations back to the primary cloud, once they're satisfied it is stable.

In the background, CloudFS will then draw on data stored locally and in the secondary cloud, incorporating all changes made while it was unavailable, to make the primary cloud consistent.



## Beyond Public Cloud Storage Outages to Enhanced Disaster Recovery

Well-publicized outages from cloud storage giants such as AWS, Azure and Google have raised awareness of the need for data availability for organizations using public cloud object storage, by ensuring that object storage itself is not a single point of failure.

While outages may be relatively short-lived when they result from factors under the cloud provider's control, the same cannot necessarily be said when and if those storage providers become the target of intentional action from bad actors.

Beyond the ability to immediately failover in the event of an primary storage outage, cloud mirroring offers you additional data durability and resilience with ransomware protection in case your primary storage provider itself is targeted with a malware or ransomware attack beyond your, or Panzura's, control.

Your secondary, mirrored store provides a real-time backup with up-to-the-moment data redundancy in case you need to switch storage providers for a time. More importantly, in the case of complete disaster, it provides accelerated recovery capabilities, by holding a redundant data set capable of granular restoration without data loss.

## **It's a Critical Layer of Cloud Data Protection**

The value your data now holds makes it imperative that it's safe, secure, durable, and available whenever it's required.

With cloud mirroring, you can provide high availability with automatic failover at the level of your chosen cloud store, guarding against cloud outages, cyber attacks made directly against cloud storage providers, cloud bucket deletion and other potentially catastrophic data storage events.

Cloud mirroring is available with CloudFS 8.1.0.