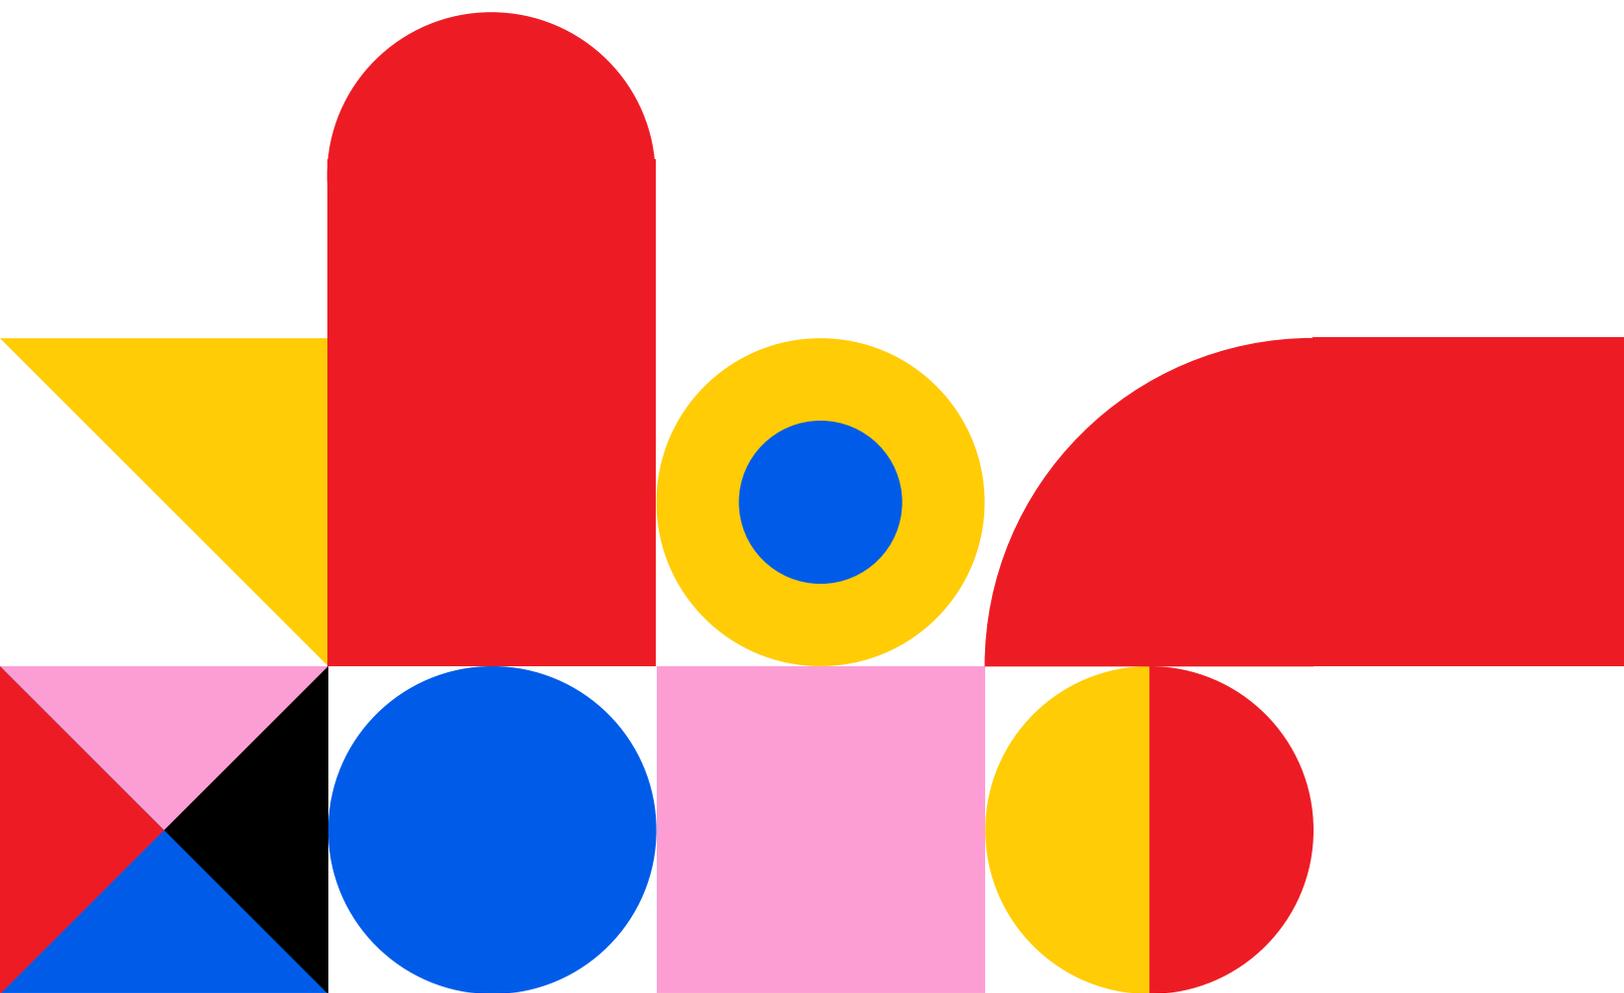


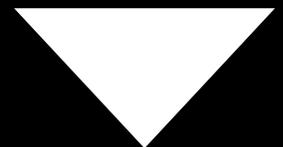
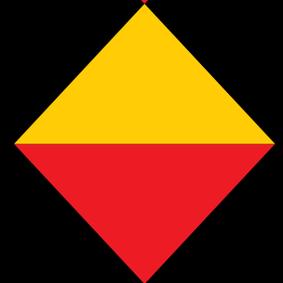
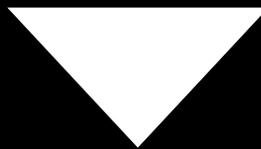
Panzura CloudFS 8

TECHNICAL WHITEPAPER



CONTENTS

- 3. Introduction
- 5. Panzura CloudFS™
- 5. File-Based Storage
- 5. Cloud Object Storage
- 6. Global Cloud File System
- 7. Global Namespace
- 8. Panzura Snapshots for Immediate Global File Consistency
- 11. Intelligent Caching at the Edge for Local-Feeling Performance
- 14. Global File Locking and Real-Time Global File Consistency
- 15. Global Deduplication
- 16. Immutable Data and Resilience to Ransomware
- 16. Military-Grade Encryption and Regulatory Compliance
- 17. Cloud Mirroring
- 18. Search, Audit, and File Network Visibility
Global Services and Customer Care





INTRODUCTION

The Network-Attached Storage that has been the data storage mainstay for nearly 3 decades is struggling to cope with escalating storage volumes. It's also unable to make files consistent across sites within a time frame that allows teams to be productive. NAS is highly performant when situated close to the users accessing the files it stores, but becomes unworkably slow when remote users attempt to access it.

As a result, organizations deploy individual NAS instances at every location, creating disconnected storage islands that contain a significant amount of data duplication. Making files consistent across locations involves scheduled data replication, if it's attempted at all, and data must be replicated to achieve an acceptable level of durability. Usually, that means at least a secondary set of data for backup, and a tertiary set for disaster recovery.

This contributes to the exponential growth of unstructured data, and creates complexity and expense for IT teams, who struggle with lack of visibility into their multiple file networks. This legacy approach to data storage has no easy way to restore granular amounts of data in case of loss, and leaves data exposed to ransomware and other malware attacks.



CloudFS transforms complex, multi-component, and often multi-vendor environments into a simplified data management solution, while addressing cost reduction, risk mitigation, and operational complexity.

Panzura's global cloud file system CloudFS provides a single authoritative data set held in cloud object storage, with immediate global data consistency and local-feeling file performance across all locations. Data from all legacy storage instances is consolidated, de-duped and compressed, significantly reducing the overall unstructured data footprint.

Compatibility with a wide array of object stores provides the flexibility to consume public cloud storage, such as AWS S3 and Azure Blob, or private object storage, such as IBM iCOS and Cloudian.

With data durability without replication*, granular ability to restore data to a point in time and resilience against ransomware, the Panzura solution not only replaces network-attached storage, but associated backup and offsite DR processes and storage. It also provides single pane of glass search and monitoring of the entire file network.

* When used with object storage that replicates data three ways.

Panzura CloudFS™

Panzura's global cloud file system CloudFS is a distributed file system incorporating network acceleration technology, specifically designed to accommodate highly latent remote object stores, and able to overcome the limitations preventing organizations from successfully integrating cloud storage into their infrastructure.

The result is a multi-cloud file services platform that enables high performance tiered NAS, global file collaboration, ransomware resilience, active archiving, backup, and DR across all an organization's locations.

File-Based Storage

Panzura developed a high-performance file-based global storage platform for the cloud to address the 80% of current data that is unstructured. By supporting NFS and SMB transfer protocols commonly used by most applications, Panzura can plug into existing IT infrastructures without any changes and connect to all major cloud storage platforms, simplifying deployment and minimizing impact on operations. All data is managed under a single global file system, simplifying user interaction and system administration while tying into organization applications and targeting both local disk and the cloud.

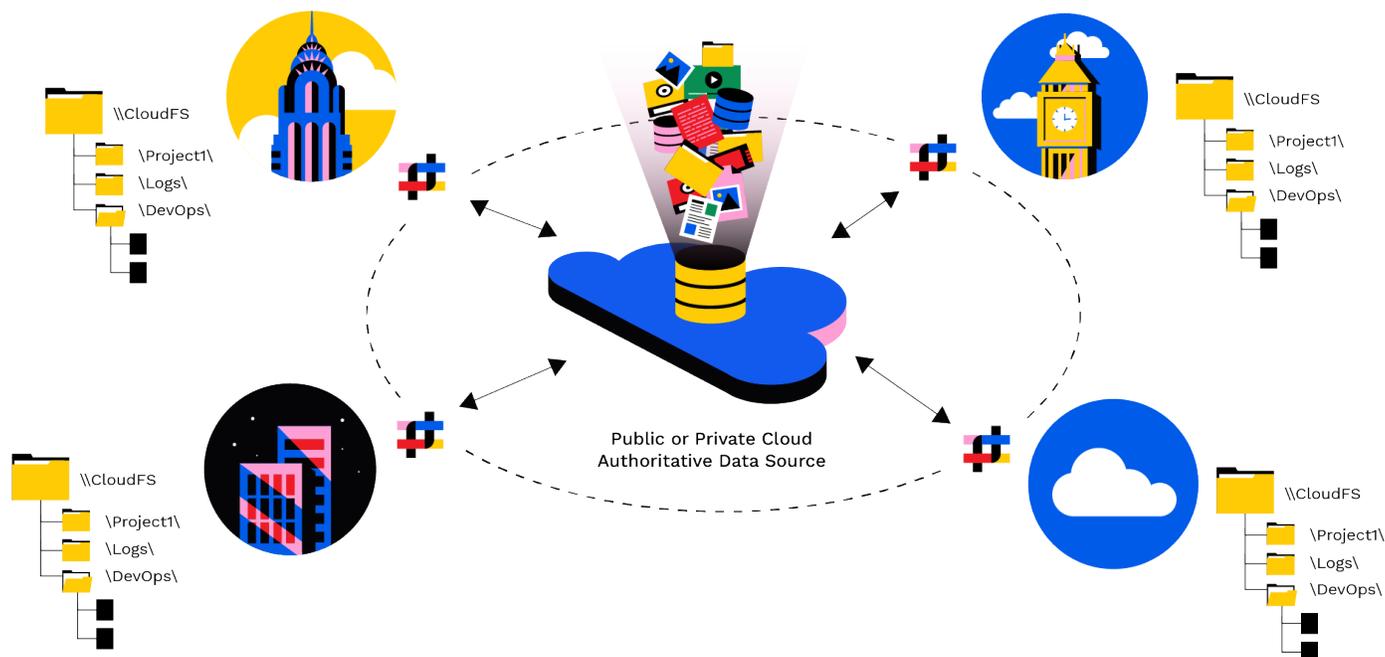
Cloud Object Storage

Object storage, the typical storage system used in the cloud, breaks up data and stores it as flexibly-sized containers or chunks. Each chunk can be individually addressed, manipulated and stored in many locations—not tied to any particular disk—with some associated metadata.

Object storage can scale to billions of objects and exabytes of capacity while protecting data with greater effectiveness than RAID. In addition, due to the discrete scale-out architecture of object storage, drive failures have little impact on data and self-healing replication functions recover very rapidly (think weeks for large capacity legacy RAID systems). This combination of scale and robustness make object storage an ideal target for warehousing data.

CloudFS interfaces directly with all major cloud object storage APIs and related storage tiers, and leverages object-based cloud storage as a data warehouse to provide scale and availability with a compelling cost structure.

Global Cloud File System



The heart of any storage system for unstructured data is the file system. Panzura CloudFS was engineered to closely manage how files are utilized and stored to provide seamless, high-performance, and robust multicloud data management. It improves on WAFL and ZFS while integrating cloud storage as a native capability. Any user, at any location, can view and access files created by anyone, anywhere, at any time.

The file system dynamically coordinates where files get stored, what gets sent to the cloud, who has edit and access rights, which files get locally cached for improved performance, and how data, metadata, and snapshots are managed. The structure of the file system has no practical limit for the number of user-managed snapshots per CloudFS.

Panzura's innovative use of metadata and snapshots for file system updates, combined with unique caching and pinning capabilities in the Panzura nodes—virtualized edge appliances deployed either locally or in the cloud—allows you to view data and interact through an enterprise-wide file system that is continually updated in real time. Support for extended file system access control lists (ACLs) empowers administrators to set file access and management policies on a per user basis.

Global Namespace

The Panzura global namespace is an in-band file system fabric that integrates multiple physical file systems into a single space and is mounted locally on each node. The entire global namespace has the root label of the distributed cloud file system.

As an example, the following 2 global namespace paths point to the same directory (\projects\team20) and are visible from both nodes as well as locally on nodes cc1-ln (London) and cc1-ny (New York).

CloudFS is designed to ensure immediate data integrity across all sites within the file system, regardless of the number of sites and how far apart they are.

This requires adherence to two fundamental principles:

1. That only one user can edit the same file – or, where applications support byte-range or element locking, the same part of a file – at any time. If another user attempts to open a file, or access part of a file that is locked for editing, they will be notified that the file is locked, or be unable to edit the file element.
2. That whenever a user opens a file with read-write access, they will see the most recent saved edits made to that file, regardless of how recently those changes were made, and the location of the user who made them.

To achieve this, Panzura decouples data from metadata, and integrates the global namespace into the metadata.

Metadata is stored centrally in the cloud for durability in addition to being fully cached locally for enhanced performance. All nodes in a single namespace or CloudFS synchronize metadata updates simultaneously every 60 seconds in a hub (cloud) and spoke (node) configuration. This is further complemented by a peer to peer (mesh) synchronization event that occurs in real-time when lock dynamically moves from one node to another through the distributed global file locking.

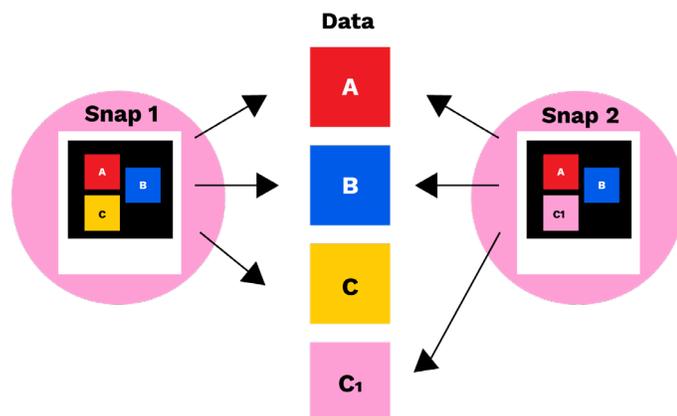
Panzura Snapshots for Immediate Global File Consistency

Snapshots for Consistency

Snapshots capture the state of a file system at a given point in time. For example, if blocks A, B, and C of a file are written and snapshot 1 is taken, that snapshot captures blocks A, B, and C to represent the file.

If someone then edits the file so that block C1 replaces C and snapshot 2 is taken, the data pointers in the snapshot file blocks A, B, and C now point to A, B, and C1. Block C is still retained but not referenced in snapshot 2. If you wanted to recover to the original state, you can restore snapshot 1, then the system will point back to A, B, and C, ignoring C1.

By using snapshots for creating and saving an ongoing series of recovery points for different stages in data's lifetime, a consistent state of the file system can always be restored in the event of a data loss or damage.

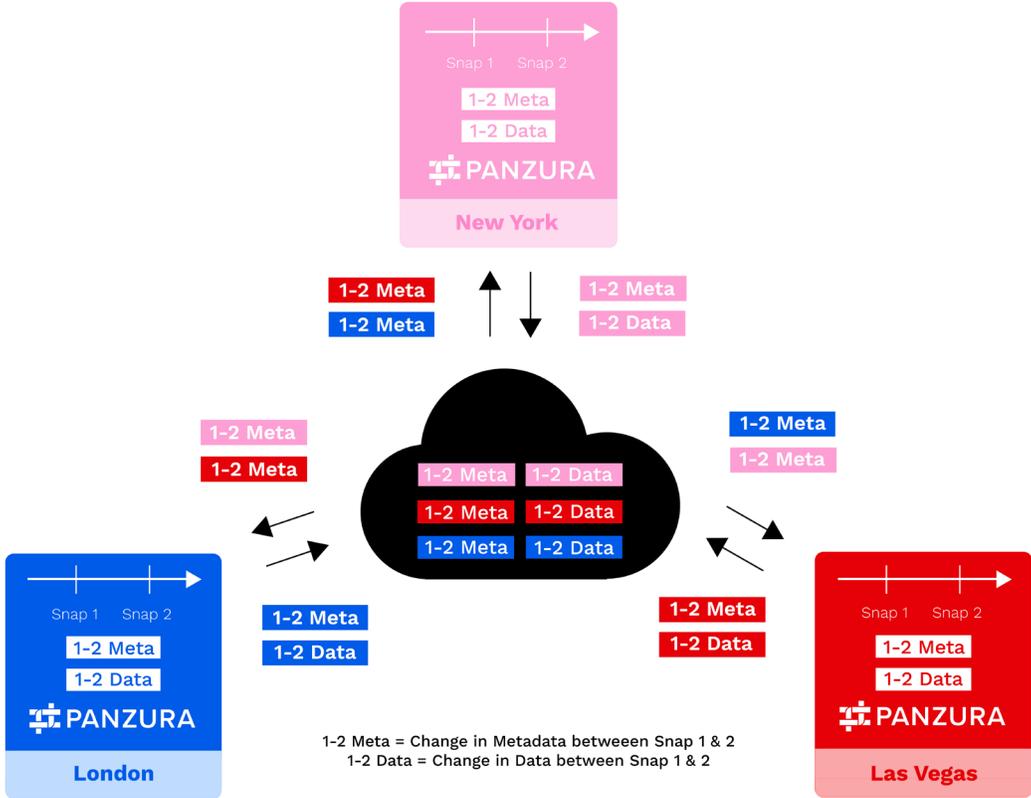


Snapshots for Currency

Panzura uses differences between consecutive snapshots both to maintain file system consistency as well as to protect data in the file system. In a process called syncing, the Panzura file system takes the net changes to metadata and data between consecutive snapshots and sends them to the cloud. The metadata portion of these changes is retrieved from the cloud by all other Panzura nodes in the configured CloudFS, where they are used to update the state of the file system and maintain currency (see image below).

This system updating occurs continuously across all nodes, with each node sending and receiving extremely small metadata snapshot deltas to and from the cloud in a hub and spoke configuration, using them to update the file system seamlessly and transparently.

For example, a node in London (blue in the figure below) takes Snap 1 and then later takes Snap 2. The difference in metadata between Snap 1 and Snap 2 for London is shown in blue as 1-2Meta. The difference in data between Snap1 and Snap2 for London is shown in blue as 1-2Data. London sends its 1-2Meta and 1-2Data to update the cloud, as do all other nodes in the infrastructure. London also receives back metadata updates for all other nodes (shown as 1-2Meta in pink for New York and in red for Las Vegas).



All of the changes in data and metadata are stored and tracked sequentially in time. Should data loss or corruption occur at the local node or in the cloud, data can be restored to any previous state at which a snapshot was taken, without the need to follow a separate backup process.

It is important to reiterate that the size of these snapshot deltas (1-2Meta, 1-2Data) are exceptionally small relative to the data in the file system; thus they can be captured continuously and use bandwidth and capacity very efficiently.

The result is the Holy Grail of a global file system: a solution that requires almost no overhead and provides near real-time, continuous rapid updates across all sites.

Snapshots for Efficiency

Panzura nodes have no practical limit for user-managed snapshots. This category of snapshots allows users to recover data without IT intervention, by simply finding the desired snapshot in their inventory and restoring it. Policies around user-managed snapshots (frequency, age, etc.) are defined by IT administration.

For example, a Microsoft Windows user in New York travels to London and realizes she needs a file that she deleted 3 months ago. She directs her Windows Explorer to the local London Panzura node, navigates to her snapshot folder, and finds the date/time that corresponds to the file system view that contains the file she wants to recover. She opens that snapshot, and navigates to the file or files she needs to recover, then just drags and drops the needed file(s) into the current file system location she wants them restored to. Within minutes, she has recovered whatever files she needs and can continue with her work, all without involving anyone from IT.

For ease of use, user snapshots have been integrated with the Windows Previous Version function allowing users to right-click on any file or folder and easily restore to any previous snapshot. IT administration can dynamically change snapshot policies as needed to satisfy data retention policies, balance frequency and duration for optimal system performance and user satisfaction.

Snapshots Benefits

Panzura snapshot technology provides three major benefits for the global file system: consistency, currency, and efficiency. Continuous snapshots provide granular recovery points so that, in the event of a data loss, a consistent file system state can be restored with minimal disruption or delay.

Panzura snapshot technology provides all users in all locations with a current view of the entire file system. This is done by syncing all file system views globally in real-time, allowing users to experience cloud storage as if it were local, solving the key inhibitor to a true global file system.

By empowering users to recover their own data as needed, Panzura snapshot technology offloads a key aspect of user support, freeing up time for strategic IT projects. The Panzura node brings the power of the cloud to organizations without sacrificing the user experience.

Intelligent Caching at the Edge for Local-Feeling Performance

SmartCache

Panzura CloudFS utilizes a user-definable percentage of the local storage as the SmartCache to intelligently track hot, warm, and cold file block structures as they are accessed. This form of caching dramatically increases the I/O performance of reads (and reduces cloud object storage access charges) by servicing them from local cached storage (both in memory and on persistent local flash) rather than from external cloud storage. The file system also buffers against variations in cloud availability to help maintain consistent read/write response times - performance AND availability at the edge.

Caching policies provide two basic functions. The first function is pinned data, which keeps data available on local storage using flexible wildcard policy rules. Pinning is a forced action and executed against full files whereas SmartCache is a read-stimulated action executed against frequently accessed blocks within a file.

Pinned data results in a 100% local read guarantee whereas SmartCache is deterministic based on previous I/O read patterns within the local node. The second function provided by caching policies is Auto-Caching which automatically caches data locally based on defined rules. However, auto-cached data can be evicted for requested hot data, as needed.

The pinned, or auto-cached, data is a subset of the total SmartCache storage tier. Pinned data is considered high-priority cached data that is never evicted unless authorized by the administrator, whereas auto-cached (cached based on wildcard rules) or SmartCache cached (data blocks automatically cached based on observed usage patterns) can be evicted by the system if needed to make space for more frequently accessed data. The balancing of pinning and SmartCache is delicate as a pinning rule will force data blocks to be logically placed inside the SmartCache, consuming local space, which may affect the local cache utilization and efficiency in ways that the administrator may not have considered. Because pinned policies are of the highest priority and override caching rules based on observed behavior, careful attention should be given to those policies so as to

not consume all of the local storage leaving little for actual hot data.

The Auto Pre-populate feature provides an even higher degree of automated caching capabilities. If enabled, the node will automatically pre-cache files based on ownership changes between nodes in a CloudFS. This is particularly helpful in collaborative workflows where users at different sites are working on the same datasets. As the node detects ownership changes between locations it will automatically cache data in the same directory in anticipation of user read requests on those files between sites.

Local Storage Usage

A portion of the local storage is allocated for SmartCache. This portion is configurable and is set to 50% by default. Over time and through general usage, the system dynamically populates the local cache with hot data blocks from all of the files being read by users and applications. The most optimal and efficient SmartCache configuration is to have most of the cache comprised of hot and warm blocks, with most cold blocks being evicted to the cloud. In this case, a high percentage of reads are serviced directly from the local cache rather than from the cloud. This is the optimal caching state, but is harder to achieve when more pinning rules are added.

Blocks residing in local cache are characterized by a combination of 3 different temperature states, 2 modification states, and 2 protection states. These are:

Pinned—Blocks that have been pinned receive the highest priority in the SmartCache and are the last to be evicted, but only if critical write space is needed.

Hot—Blocks frequently being accessed for reads (from 0-7 days).

The goal is to have mostly hot blocks in the local cache.

Warm—Blocks that were recently hot but have not been read as recently as any of the hot blocks (8-30 days). They will be evicted after cold but before any hot blocks if extra SmartCache space is needed.

Cold—Blocks that have not been accessed for 30 days or more. These are the first blocks to be evicted when SmartCache needs space for pinned, hot, or warm blocks. There should always be some cold blocks as this indicates that the SmartCache completely holds all pinned, hot, and warm blocks.

Recently modified—Blocks that have been written to as part of updates to a file.

Not modified—Blocks that have not been written.

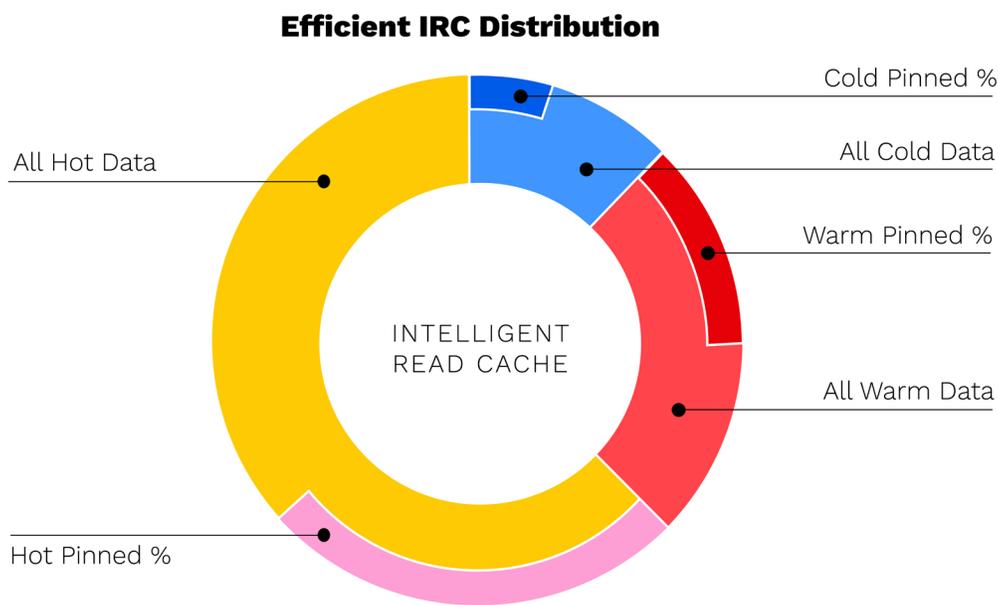
Protected in the cloud—Blocks that have been successfully uploaded to the cloud storage.

Not yet cloud protected—Blocks that are pending upload the cloud.

Pinning consumes SmartCache space by forcing complete files into the local cache and is designed for the administrator to satisfy user or site needs by overriding the SmartCache's auto-caching logic to disable eviction indefinitely for specific blocks. Because of this, careful attention to specific pinned rules should be given to prevent a rule that could cause thrashing of the local cache space (rotating eviction of data with new data due to reduced cache capacity). It is recommended that administrators utilize the Auto-Caching action or enable the Auto Pre-populate feature where possible. Panzura CloudFS is designed to transition all data into the cloud as quickly as possible. Data is always committed and uploaded to the cloud before becoming hot, warm, or cold based on any recent read activity.

When data is pinned, that data is only evicted from SmartCache if the administrator changes the pinning policy or space is needed for writes and all other hot, warm, and cold data has been evicted.

Pinned data is considered high-priority cache data. Inversely, auto-cached SmartCache data is treated as lowpriority cache data that can be evicted automatically by the system as SmartCache space is needed for new hot data. As more pinned data consumes the IRC, the usable auto-cache capacity is reduced. This will negatively impact the most frequently read data, causing it to be evicted and then re-read continuously. Therefore, aggressive policies that pin large amounts of data should be used sparingly as this could cause excessive local disk I/O and reduce performance.



Ideally, most of the data that applications need should be resident in the local cache. The diagram at right depicts a case where all hot and warm data is auto-cached with some cold data and some pinned data. Overall, most of the Smart Cache local-disk space is being used by active data (hot+warm). The amount of cold pinned files should always be monitored as this indicates a pinning rule that is no longer relevant and potentially no longer needed. Those rules should be removed from the system.

Global File Locking and Real-Time Global File Consistency

Panzura is the only global file system with real-time data consistency across all sites. That is, any user opening a file for editing will see the most recent saved changes, regardless of where those changes were made. Our patented file locking process plays a crucial role in this process. Global file locking is at the heart of allowing geographically distributed users to work collaboratively, without overwriting each other or creating multiple file versions.

Data Ownership, Data Locking and Data Mobility

CloudFS physically decouples data and metadata. This decoupling enables the file system to be highly flexible in referencing which physical blocks are used to construct a file. It also allows every node in the file system to hold a complete copy of the metadata for the whole file system, without having to hold the files themselves.

Panzura's global distributed file locking follows three simple principles.

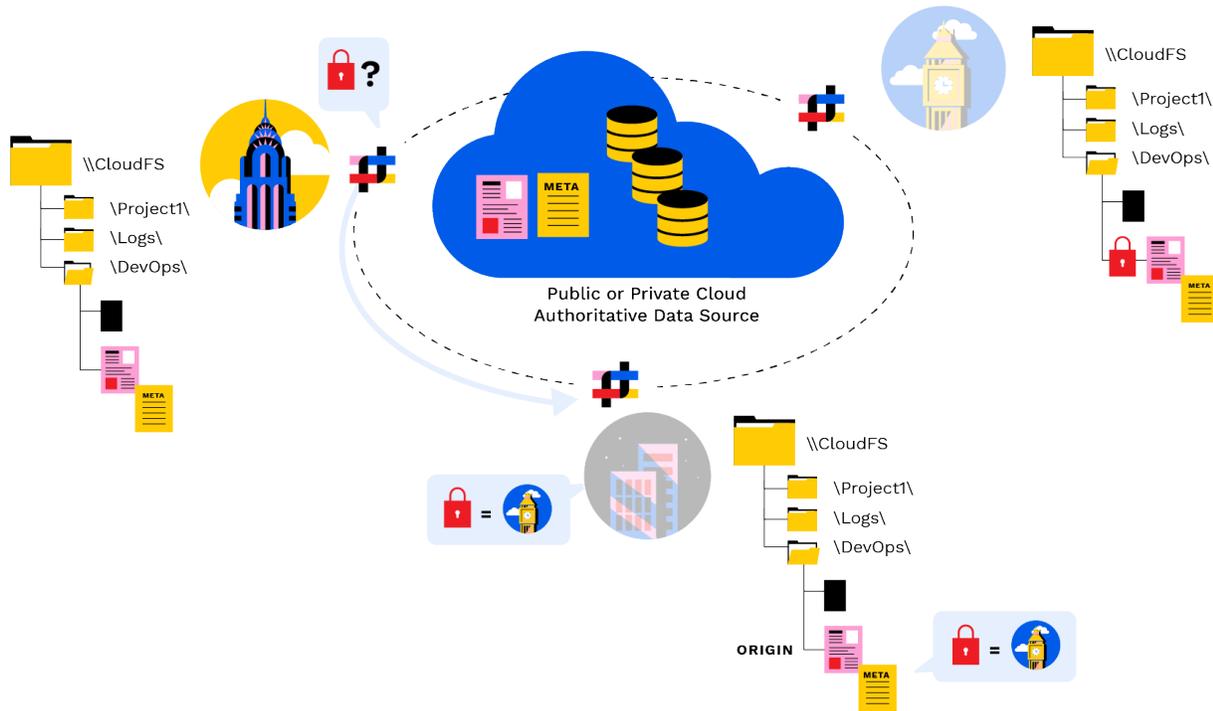
1. When a file is created, the node on which it was created is designated the Origin, and this is recorded in its metadata.
2. The Origin always knows which node currently has the lock, regardless of whether the file is currently locked for editing.
3. The node with the lock is the Data Owner, and this information is held in the file's metadata.

The Data Owner state is transported via metadata snapshots. A node wanting to assume Data Ownership for a file checks its metadata for the node on which the file was created (the Origin) and then communicates directly with the Origin, to request the lock and become the Authoritative Write Node.

If the lock is sitting with the Origin, it will either approve the request, or deny it if the file is open for editing. If the Origin does not currently hold the lock, it tells the requesting node which location to request it from.

Data Ownership requests and transitions are frequent events and are negotiated in real time via small peer-to-peer communications among nodes.

The final step after a Data Owner transition is to ensure the user now opening the file sees any changes that have been made to the file since the last sync to the object store. This involves a direct peer-to-peer communication between the Origin and the new Data Owner, and possibly the current Data Owner (which might not be the Origin).



Within this peer-to-peer stream, the ownership metadata computes a final delta list of real-time changes that may have occurred since the Data Owner changed.

This list, which can be as small as a single file system block, is streamed directly to the new Data Owner via a secure optimized data channel. The new Data Owner processes all remaining deltas, making the file current and consistent.

All file reads and writes from that Panzura system now happen as local I/O operations on the new Data Owner. The Data Owner retains full read/write ownership until a new Data Owner transition occurs.

Global Deduplication

Panzura's interconnected global file system stops file-level duplication before data gets synced to the object store. Since only unique copies of files across all sites are preserved by the file system, data is deduplicated before it is ever stored.

Capacity is optimized further by running advanced, inline block-level deduplication on any data in the object store, in order to remove blocks common across different files. Unlike any other deduplication provider, Panzura embeds the deduplication reference table in metadata, which is instantly shared among all Panzura nodes. This inline deduplication method removes data redundancy across all nodes, rather than just based on data seen by a single node. Thus, each node in the network benefits from data seen by all other nodes, ensuring even greater capacity reduction, guaranteeing all data in the cloud is unique, driving down cloud storage and network capacity (and cost) consumed by the enterprise.

Immutable Data and Resilience to Ransomware

The persistence, pervasiveness, and documented success of ransomware attacks would suggest it may not be possible to mount a complete first-line defense, even within well-resourced organizations. That makes it essential that critical business data is as close to invulnerable as it can possibly be. That is, if your environment is attacked, and even accessed, the data itself will not fall.

At the heart of every ransomware attack is the ability to encrypt files such that they cannot be accessed or recovered without paying a ransom to the attackers, in return for the ability to decrypt them. Panzura makes data impervious to ransomware by storing it in an immutable form (Write Once, Read Many) and further protecting it with read-only snapshots.

With Panzura, once data is in the cloud object store, it cannot be changed, overwritten, or damaged in any way. File changes are written as new data blocks, which have no effect on existing data. As new data is saved, Panzura's global file system updates file pointers to record which data blocks comprise a file at any given time.

Panzura's lightweight, read-only snapshots then provide a granular, point-in-time ability to recover any data, by restoring from the applicable snapshot. Individual files, folders, or even the entire file system can be restored in this way.

Because both the snapshots and the data itself are immutable, ransomware attacks do not damage files in the Panzura global file system. Instead, attacks are shrugged off by quickly reverting back to previous data blocks, to make up uninfected files.

Military-Grade Encryption and Regulatory Compliance

Panzura addresses data security concerns directly by applying military-grade encryption to all data stored in the cloud. Each Panzura node applies AES-256-CBC encryption for all data at rest in the object store. In addition, all data transmitted to or from the cloud is encrypted with TLS v1.2 while in flight, to prevent access via interception. Encryption keys are managed by the organization, never stored in the cloud. The solution is FIPS 140-2 certified.

This complete, robust two-tier encryption solution is in addition to the typical multi-layer security provided by mainstream cloud storage providers. In some cases, companies find that the combined security of a Panzura+cloud solution is greater than they can reasonably achieve within their own infrastructure, making cloud storage safer than some private cloud deployments.

Secure Erase

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. All data is securely erased and replaced with zeros. Secure erase can be used with any supported cloud provider.

Cloud Mirroring

Using cloud mirroring, you can effectively double the availability SLA of any single cloud storage provider while providing uninterrupted service in the case of a cloud storage service outage. Cloud mirroring will automatically failover to a redundant cloud storage provider in the case of a failure of the primary provider **without disrupting any front-end file services for systems or users.**

This is only made possible because cloud mirroring delivers immediate data consistency. Failover at time of failure is not possible with eventual data consistency, which is what most other replication features offer. When the primary cloud object store is back up, Panzura will automatically synchronize both clouds to a consistent state—all without human intervention. Additionally, you are protected against accidental object or bucket deletion.

The cloud mirroring functionality addresses problems of auto-failover in case of cloud failure, provides a full backup beyond single cloud replication and automatically initiates syncing of clouds after failure. As organizations increasingly employ multiple clouds for storage, cloud mirroring helps by eliminating dependency on any one vendor.

Search, Audit, and File Network Visibility

Panzura's powerful SaaS data management solution Data Services provides a single, unified view and management of unstructured data, whether it's stored in the cloud, on premises, or at the edge. Data Services strips hours out of daily IT administrator activity, as well as being a valuable tool for rapid recovery from ransomware attack.

Global Search

Accelerated global search finds files in seconds, searching across your Panzura CloudFS and any other connected nodes. From search results, audit and file recovery options are available with one click.

File Audit

Files can be queried by user action, as well as by user. Using audit actions such as renaming files or setting file attributes can narrow a search to find potential data damaging actions that may contain ransomware, while actions such as open, and copy can pinpoint potential unauthorized access of data.

Clone and Replace

Can revert damaged or deleted files to previous versions, and to previous locations, in seconds.

File Analytics

Storage metrics at a glance assist administrators to understand what's consuming space, how storage requirements are changing, what's most frequently accessed, and which users are most active.

File System Pulse

Proactively monitor file system health metrics such as CPU usage, data movement, cloud connectivity, and more.

Global Services and Customer Care

Panzura provides service and support at every stage in your Panzura deployment, from solution design and build, to ongoing operations and support.

Technical support is available 24/7/365, and cost-effective support options allow you to determine the service level you require, from rapid reactive support to proactive and predictive support designed to keep missioncritical services running at maximum efficiency.

Panzura has an industry-leading NPS score of 87.



SUMMARY

Replacing legacy storage with a modern approach to unstructured data, using cloud object storage, offers tremendous potential for organizations to reduce storage costs, improve productivity, and reduce data availability risk.

Tapping that potential fully and effectively can provide significant competitive advantage while reducing both business and technological risk. Critically, Panzura empowers organizations with immunity to ransomware, allowing rapid recovery in the event of attack, minimizing loss of data, time or productivity and removing the need to pay a ransom.

The unique business value delivered by Panzura includes:

Radical reduction in nodes/appliances and overall complexity

Panzura eliminates multiple workloads and disparate file systems. All solutions are delivered from a single vendor, lowering support costs and required in-house expertise.

Built-in reductions to current risk profile by more than 75%

The immutable data storage architecture and replicated cloud object storage eliminates the need for additional BC/DR/Backup solutions while adding ransomware protection, encryption at rest, and data lifecycle management.

Evergreen product that enables your ideal future state

The software-defined solution eliminates hardware refreshes and migrations while serving as the same transport layer to the object front end you are moving to in the future.

Unlocks data for next generation workloads for a competitive advantage

Accelerate your time-to-value from AI/ML/NLP using Azure Cognitive services or other cloud APIs by having the full dataset available everywhere.

Achieve true zettabyte-scalable file systems

This architecture eliminates many of the bottlenecks of traditional file systems where objects are constrained by inodes, pathnames, or address allocation.