PANZURA

PANZURA CLOUDFS SOLUTION BRIEF

Banking, Financial Services, and Insurance

The financial services industry is undergoing a massive transformation, driven by the need to deliver personalized customer experiences, optimize risk management, and comply with increasingly stringent regulations. This transformation is fueled by an explosion of data, including customer transaction histories, market data feeds, regulatory filings, and complex financial models. Effectively managing and leveraging this data is no longer just an operational necessity, but a critical factor in maintaining a competitive edge and ensuring long-term stability.

The sheer velocity and volume of data generated—including high-frequency trading data, massive customer databases, and complex actuarial models—are overwhelming traditional storage infrastructures, hindering timely analysis and decision-making.

When a risk analyst needs immediate access to real-time market data to assess potential exposures, or an underwriter requires seamless collaboration on large policy documents across multiple branches, the performance and reliability of the underlying data infrastructure become paramount. The burden on IT to ensure this while maintaining stringent security and compliance is immense.

Navigating the complex landscape of regulatory compliance, including GLBA, SOX, GDPR, and other regional mandates for protecting sensitive financial and customer data, while simultaneously defending against increasingly sophisticated cyber threats targeting valuable financial assets and customer information, creates a critical need for robust data governance and security.

Whether embracing the scalability of secure public clouds, implementing private cloud solutions, or adopting a hybrid approach, a comprehensive data management strategy is essential to ensure secure access, robust protection, and seamless collaboration across the financial services ecosystem.

In this solution brief, we outline how Panzura helps banking, financial services, and insurance organizations address the complex and evolving challenges related to data storage, protection, access, and management. The escalating data deluge, coupled with stringent regulatory demands, persistent cybersecurity threats, the imperative for seamless data interoperability across disparate systems, and the growing reliance on advanced analytics and AI for fraud detection and algorithmic trading, highlights the indispensable need for sophisticated, scalable, and secure data management solutions like Panzura CloudFS.

Consolidating Distributed Data

Data distributed across multiple storage devices, and often across multiple locations, is prone to a significant amount of duplication as similar and sometimes almost identical files occupy storage space. Myriad similar copies of the same file cause confusion among users about which version they should rely on as the authoritative copy, making compliance difficult.

While cloud storage presents with a cost-effective pricing structure, simply lifting and shifting that data into the cloud replaces one set of data islands with another. Performance problems working with cloud-stored data, coupled with storage inefficiencies and lack of visibility can quickly render a cloud move ineffective. Panzura CloudFS consolidates distributed data into a single, authoritative data set that is visible, and accessible, across the organization.

Minimizing Duplication

CloudFS deduplicates redundant data before moving it to your chosen cloud or object store, allowing you to realize a significant reduction in your overall data footprint. CloudFS maintains this globally deduplicated data set at all times, checking for redundancies every time it moves data into your cloud storage.

Accessing Data in Real Time

With CloudFS, all users in your organization work from the authoritative data set stored in your cloud or object store. No changes in workflows, or user behaviors are required – users interact with files in the same way they always have, and CloudFS provides them with a local-feeling file experience. That means that files open and save as quickly as they always did when stored locally.

Real-Time Data Consistency for Critical Financial Operations

CloudFS minimizes risk in banking and financial operations by ensuring that file edits are immediately visible everywhere upon saving. This guarantees that even time-sensitive data, including financial models and trading information within Excel files dependent on macros, which can slow file operations over distance to a crawl, is instantly consistent across all locations. This immediate file consistency means users can always access the authoritative file, with the most recent changes, eliminating delays and inconsistencies that can lead to costly errors and missed opportunities.

Seamlessly Move Critical Applications to the Cloud

With Panzura, there's no need to rewrite legacy applications to support your move to the cloud. CloudFS allows file-based applications to read and write data to the cloud, allowing you to move even your most critical workflows to data using cloud storage, without changing applications, or the way your users work.

Empowering Collaboration

For workloads that require people to work together on the same files, without running over the top of each other, CloudFS empowers real time collaboration in a way nobody else can. Instantaneous, automatic file locking locks down a file for editing the moment it's opened. Where applications support element or byte-range locking, such as Microsoft Excel, CloudFS allows multiple users to work within the same file, without overwriting each other.

Keeping Data Protected

CloudFS provides inherent protection against accidental file deletion or corruption stemming from malware or ransomware attacks through its resilient data architecture. All data managed by CloudFS is stored in an immutable – Write Once, Read Many (WORM) – format. Once a data block is committed to your object storage, it cannot be altered or overwritten.

Every new edit or newly created file is stored as new data blocks, which never overwrite existing blocks. Consequently, in the event of a malware or ransomware incident, the existing, pristine data remains untouched and recoverable, as the malicious software cannot modify the immutable data.

Catch and Stop Ransomware

An extended capability of the CloudFS hybrid cloud platform, Panzura Detect and Rescue identifies ransomware in real time and stops it automatically by switching off the affected users, followed by a comprehensive ransomware tracker to help administrators rapidly identify and recover damaged files. Meanwhile, CloudFS's data insights and intelligence layer — Panzura Data Services — enables configurable alerting on suspicious user behavior, e.g. multiple file copy or move actions that may indicate data exfiltration.

Restore Damaged or Lost Data

In the event of any file damage – whether caused accidentally or as part of a wider encryption attack such as a ransomware event – individual files, folders, or the entire file system can be restored to a pristine state with no data loss, and minimal disruption.

Read-only global system snapshots are taken on a scheduled basis, and record the file system at that point in time. Additionally, snapshots are taken at every site location in the CloudFS every 60 seconds. This provides the ability to restore any file to any point in time as required, with a global recovery point that is never more than 60 seconds.

Ensuring Data Security and Regulatory Compliance

CloudFS empowers your institution with Cyberstorage, seamlessly integrating NIST Cybersecurity Framework functions to establish a robust, multi-layered defense for your sensitive financial data. By leveraging Cyberstorage, you can ensure your organization adheres to stringent regulatory requirements and maintains the trust of your clients and stakeholders throughout the data lifecycle.

CloudFS's built-in end-to-end encryption, immutable storage, and granular access controls rigorously protect your critical financial assets—including transaction records, customer data, trading information, and proprietary financial models—against unauthorized access, data breaches, and sophisticated cyber threats such as ransomware.

This architecture helps meet the requirements of key regulations. For example, CloudFS helps financial institutions comply with the Gramm-Leach-Bliley Act (GLBA) by ensuring the security and confidentiality of customers' nonpublic personal information. It also aids in compliance with the Sarbanes-Oxley Act (SOX) by mandating the protection of financial records and internal controls.

Additionally, FINRA, the Financial Industry Regulatory Authority, has requirements regarding data preservation and protection, especially for broker-dealers, and CloudFS's immutability ensures records are stored in a non-erasable and non-rewritable format, which is crucial for regulatory compliance and audit trails. For financial institutions that process credit card data, CloudFS's security features help protect cardholder data and maintain the integrity of transactions, thereby aiding in Payment Card Industry Data Security Standard (PCI DSS) compliance.

FIPS 140-3 certification ensures that your data remains securely encrypted both in transit and at rest, rendering it unreadable even if intercepted. Coupled with strict, role-based access controls, comprehensive audit logs, continuous monitoring for suspicious activity, and automated compliance tracking, CloudFS facilitates adherence to these regulations and mitigates the risk of non-compliance.

This proactive approach safeguards your institution's financial stability, protects valuable assets, and preserves client trust.

Empowering High Availability

CloudFS meets strict requirements for highly resilient, highly available file services. Every location in a global file system always has read access to data from every other location. Data is stored securely in the cloud and each location can read that data. In the event of a disaster in one location, every other location already has access to the data for immediate recovery.

Three options for CloudFS virtual nodes offer high availability to suit your requirements and budget.

1. Local high availability uses an active/passive stand-by pair of nodes that offer rapid failover.
2. With global high availability, in the case of a regional outage, a stand-by node will assume lock management for the failed CloudFS node.
3. Instant Node offers a sub-5 minute recovery, inclusive of boot time, with no dedicated stand-by node required. Instead, Instant Node utilizes available virtual machine CPU and memory.

Cloud Mirroring provides high availability for your object store by enabling a passive, identical copy of your data in a secondary hyperscaler or low cost object store provider such as Wasabi, Backblaze or Seagate Lyve Cloud. In the event of a primary object store outage, all CloudFS nodes will fail-over to the secondary store for read and write operations, with no disruption to users.

Regional Store allows globally dispersed organizations to operate up to 4 active copies of the object store in different cloud regions offered by their choice of AWS or Azure. These regional buckets are synced via the hyperscaler back-end network and allow office locations in each region to read and write data over the shortest possible distance to maximize performance. Should a single object store become unavailable, the CloudFS nodes will fail over to an object store in the next closest region.

Ready for AI

Training Large Language Models (LLMs) on unstructured data enables improved fraud detection, more efficient regulatory compliance, enhanced customer support experiences, accurate market sentiment analysis, and streamlined insurance claims processing. Additionally, LLMs facilitate robust know your customer and anti-money laundering processes, faster contract analysis, better-informed credit underwriting, efficient internal knowledge management, and stronger cybersecurity defenses.

Leveraging LLMs with unstructured data allows financial services organizations to proactively manage risks, reduce costs, and enhance operational efficiency and customer satisfaction.

Accelerate AI training on unstructured data by strategically positioning your data closer to your Large Language Models (LLMs) with CloudFS. If you're using a cloud-based AI solution hosted by the same provider as your object storage, simply deploy a CloudFS node within the same cloud region.

This setup allows you to leverage CloudFS's S3 interface to efficiently scan your file system without incurring egress fees. For on-premises LLM deployments with cloud-based object storage, be aware of potential egress charges, or opt for an object store provider that offers zero egress charges.

To further enhance security and efficiency for advanced AI use cases, explore how Panzura Symphony functions as a Zero Trust Data Broker, streamlining secure and compliant data access for training your AI models on sensitive content.

Work with your data, the way that works for you

Every part of Panzura's data management solution has been specifically and intentionally designed to let you put data at the fingertips of the people who need it, the moment they need it, while keeping it secure, protected against threats, and compliant with external regulations as well as internal mandates.

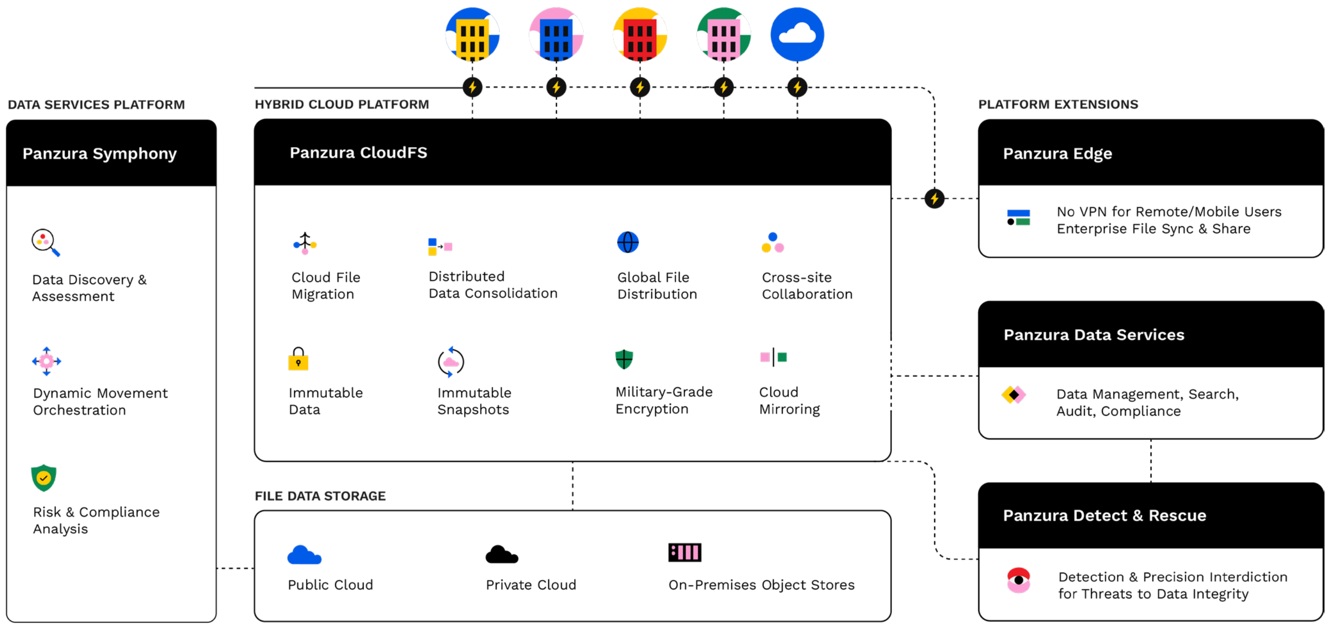
DATA SERVICES PLATFORM
Panzura Symphony
Data Discovery & Assessment
Dynamic Movement Orchestration
Risk & Compliance Analysis
HYBRID CLOUD PLATFORM
Panzura CloudFS
Cloud File Migration
Distributed Data Consolidation
Global File Distribution
Cross-site Collaboration
Immutable Data
Immutable Snapshots
Military-Grade Encryption
Cloud Mirroring
FILE DATA STORAGE
Public Cloud
Private Cloud
On-Premises Object Stores
PLATFORM EXTENSIONS
Panzura Edge
No VPN for Remote/Mobile Users
Enterprise File Sync & Share
Panzura Data Services
Data Management, Search, Audit, Compliance
Panzura Detect & Rescue
Detection & Precision Interdiction for Threats to Data Integrity

Panzura empowers today's digital-first organizations to do impossible things with file data, making them more agile, efficient, and productive. They trust Panzura to help them consolidate dispersed data, see and manage data in and out of the cloud, make it more cyber-resilient and AI-ready, and ensure it is available to people and processes where and when it's needed.

Discover how Panzura can fuel your success at **panzura.com.**