PANZURA

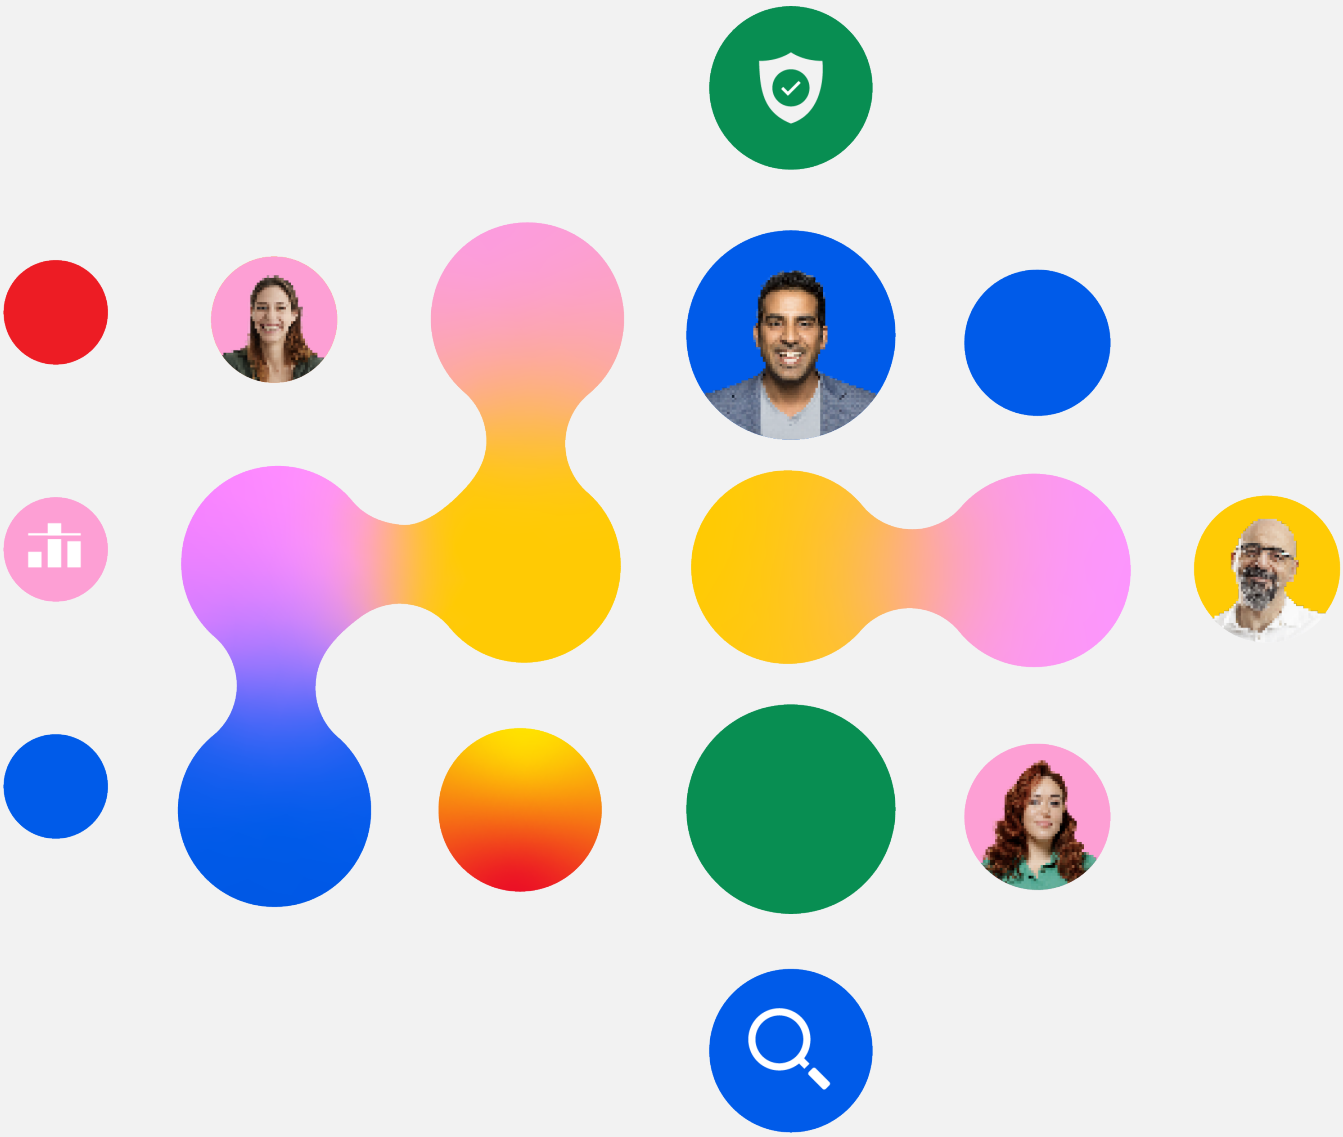

WHITEPAPER

Driving Forward in the Digital Era

A Blueprint for Resilience and Advantage

Table of Contents

Digital transformation has ushered in limitless opportunity and immense complexity in equal measures. Data was important, rapidly became essential and now, organizations are totally data-dependent. As a result, data is both more valuable and more vulnerable than it has ever been.

Against this backdrop, there are three serious threats to organizational stability and competitiveness.

01 Cyber attacks are inevitable and the stats suggest that it may not be possible to create a completely impenetrable data estate. Business leaders are faced with the reality that their defenses will be breached and their job, and the survival of the organization will depend on their ability to quickly recover data and maintain business continuity.

02 Data volumes continue to spiral out of control and lack of visibility makes it almost impossible for IT and operational teams to reliably understand what they are storing, how it should be treated, what value it holds and what risks it presents to the business.

03 Organizations are increasingly distributed across multiple locations and that makes it extremely difficult for users, workflows, and processes to access the same dataset at a speed that supports the organization. Instead, IT teams resort to file replication and boost transfer speed with expensive acceleration techniques, or leave users with file access that slows them down. As a result, organizations struggle to consistently get the best work from their best people because they simply can't get to files in a way that allows them to collaborate effectively.

The dilemma that IT and storage teams face is how to accommodate a move to cloud object storage, boost data resilience, and achieve data compliance without compromising performance and productivity. If that weren't challenging enough, they also need to find a way to control the spiraling storage costs that result from exponential increases in file data volumes and subsequent data backup requirements.

In this whitepaper, we explore the framework that empowers organizations to overcome these challenges and position themselves for resilience against threats, command and control of data, and immediate data delivery to gain powerful competitive advantage.

PART ONE

Shrugging Off Data Disasters: Building Resilience for Business Continuity

Digital transformation has introduced incredible opportunity in all aspects of our personal and business lives. Organizations can now scale in ways that were previously unthinkable, as they are no longer tethered to any location or restricted by processes that require human input at every step.

It has also caused an exponential expansion of threats. It's no exaggeration to say that IT leaders are responsible for managing a level of risk that has never been higher than it is right now.

The myriad interconnected systems that underpin business operations and the sheer volume of data that's being generated also create a level of complexity that is extraordinarily difficult to manage. Not only are applications, systems and data spread across multiple on premises and cloud locations but users — who work with and generate even more data — can be located anywhere and access systems using devices that are difficult to completely secure.

The multiple tools and solutions deployed to manage and monitor mission-critical operations and data pile on yet another layer of complication for IT teams that already feel overwhelmed.

The inevitable result is that organizations tend to overestimate their actual level of resilience. Now, all that it takes to bring an organization to its knees is to lose access to data. Survival is entirely dependent on its ability to restore pristine data quickly enough that business operations suffer minimal disruption. Available reports on the impact of cyber-attacks and the extended worldwide outages experienced as a result of a failed Crowdstrike update clearly indicate a disconnect between theoretical levels of business continuity and real-world responses to unforeseen business disruption.

In the realm of managing file data in the cloud, there are two crucial elements that IT leaders must allow for. The first is the ability to quickly and easily restore pristine data when it is damaged. The second is to maintain continuity in the event of any storage outage.

Leveraging Immutability for Active Data

Immutability, within the context of data storage, refers to the concept that once data is written, it cannot be modified or erased. Immutable storage media has long been a solution for data that doesn't need to be accessed frequently. For example, tape and disk backups preserve data in an immutable form but are slow and cumbersome for data restoration. When properly configured by the provider, object storage can also provide data immutability. For this reason, it is already widely adopted for data archival or backup purposes.

However, adoption of object storage for the most active data is less widespread. Yet active file data is often among the most valuable data an organization owns. It is also the most vulnerable.

Alongside a solution capable of overcoming performance overhead and increased data management complexity, object storage offers powerful active data benefits. Several forces have combined to intensify focus on leveraging the immutable capabilities of object storage. The first is the move to the cloud itself. While it's difficult to identify a precise figure, object storage constitutes a significant and growing share of cloud storage. IDC estimates that unstructured data, which is typically stored as objects, represents as much as 80% of all data worldwide. It is evident that, as enterprises continue to transition workloads to the cloud, files present a particular challenge due to their unstructured nature.

Enterprises face extraordinary challenges in the threat landscape, which is primarily associated with cybercriminals. Every day, they must defend against efforts to breach their systems and compromise the integrity of their data. Implementing a Zero Trust security model goes hand-in-hand with immutable storage paradigms. Data can be compromised through various means, not just malware encryption, including human error or intentional deletion.

Additionally, there is a growing demand for regulatory compliance, especially in industries where every modification to a file must be recorded in a way that cannot be altered, ensuring it can be accessed when needed.

Panzura's hybrid cloud platform is designed for the modern cloud era and it has immutability at its core. The backbone of the platform is CloudFS — a global cloud file system that ensures consistent performance, security, and data immutability whether deployed on-premises or in the cloud. To users, CloudFS looks and behaves like a standard Windows file share. However, behind the scenes is a radically different approach to moving, storing, and editing data that imbues organizations with resilience that allows them to maintain file operations even in the face of a widely damaging event.

The immutability feature of CloudFS is especially important and has become more prominent in today's cybersecurity environment. By incorporating immutability into the cloud-native architecture of Panzura CloudFS, the file data stored in the system is not just secure but also able to withstand constantly changing data risks.

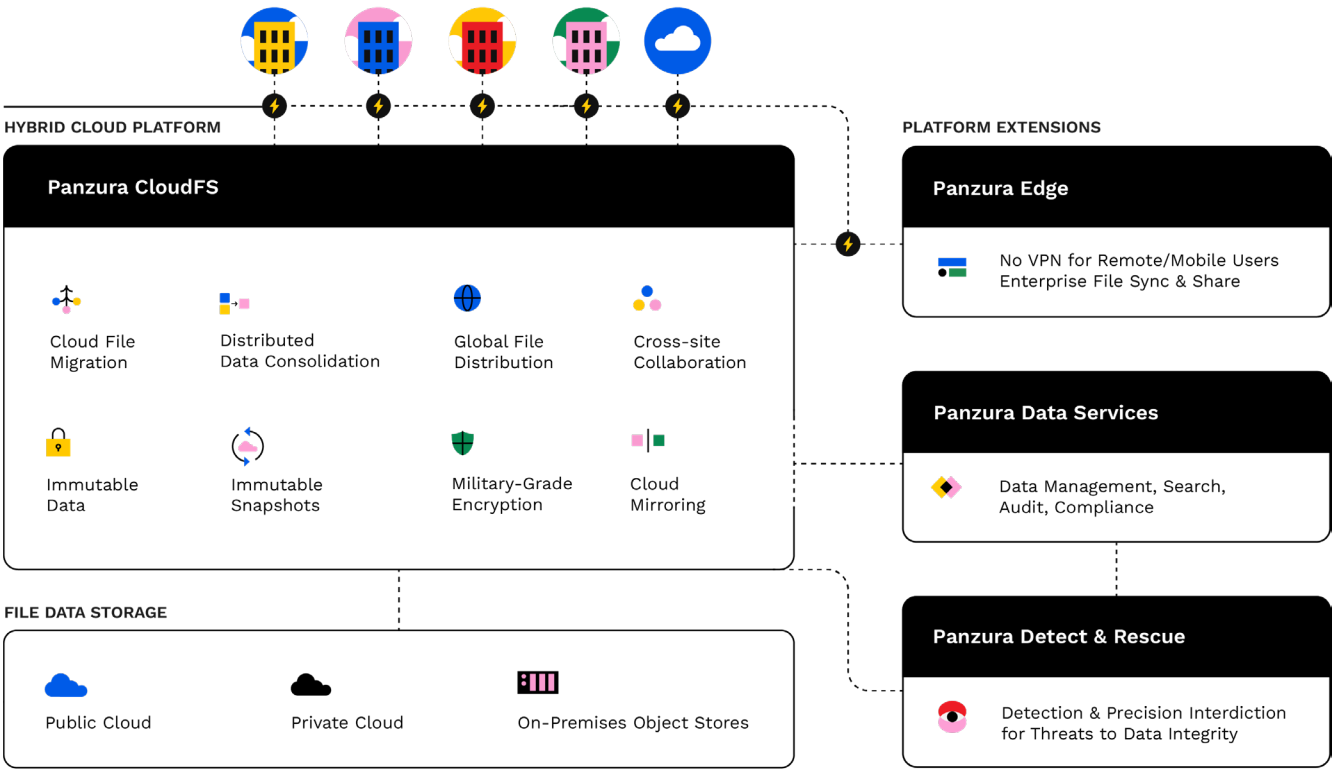

This focus on immutability is not an afterthought, but a core component that helps organizations meet regulatory requirements and other IT guidelines. Immutable snapshots offer precise restoration points, which means CloudFS can protect data from tampering and cyber-attacks such as ransomware, and it can enforce data retention policies. This gives organizations confidence in the veracity and integrity of their data in the long term.

Immutability as Protection for Editable Data

Implementing immutability in file systems poses a distinct set of technological hurdles. Conventional file systems are not designed with immutability in mind. In fact, they are explicitly architected to allow files to be altered and this leaves organizations dependent on backups if data is edited in ways that damage it.

Panzura has successfully navigated the technological challenges of implementing immutability in file systems by focusing on purpose-built architecture, performance optimization, and customizable policies. We have demonstrated that with the right approach, the benefits of immutable storage can be realized without compromising functionality or the user experience.

This has involved creating an architecture that inherently supports immutability, rather than attempting to modify existing systems. It's a design choice that undergirds CloudFS, allowing for more seamless integration of immutability features and providing a solid foundation for secure data storage.

The security measures required for immutability can potentially slow down file access and transfer speeds. Panzura CloudFS addresses this by optimizing for performance to achieve the "impossible" — real-time global file data consistency for all users. We achieve this through proprietary, patented advanced caching techniques and efficient data management algorithms that minimize latency and maximize throughput by splitting metadata from file data and moving the least amount of data as fast as possible.

Balancing immutability with agility becomes even more challenging as modern data environments become more complex, with a wide array of data sources, the need for real-time analytics, and the rise of cloud native architectures. The difficulties are exacerbated by the growing demand for enhanced data processing capabilities, stringent data governance and security standards, and the integration of artificial intelligence (AI) and machine learning (ML) technologies into data infrastructure.

Immutability does not mean that data cannot be deleted in accordance with retention and data hygiene policies. CloudFS offers systematic support for deleting immutable data, with an audit trail, when required. When compared to other global file system solutions, this approach stands out.

In addition to the established benefits of Panzura CloudFS over conventional NAS storage in terms of cost and resiliency, a more detailed examination reveals important distinctions that greatly affect practical, real-world performance.

For example, Panzura provides a single authoritative data set held in cloud object storage, with immediate global data consistency and local-feeling file performance across all locations. This greatly reduces the overall unstructured data footprint, while ensuring data durability without replication.

When hit by ransomware, numerous organizations have confronted the fact that their stated service level agreement (SLA) for recovery of clean and usable data has not survived contact with the real world. Data restoration from backups can be a lengthy process due to the size of the datasets and the process can inadvertently overwrite recent changes in files that had not been damaged.

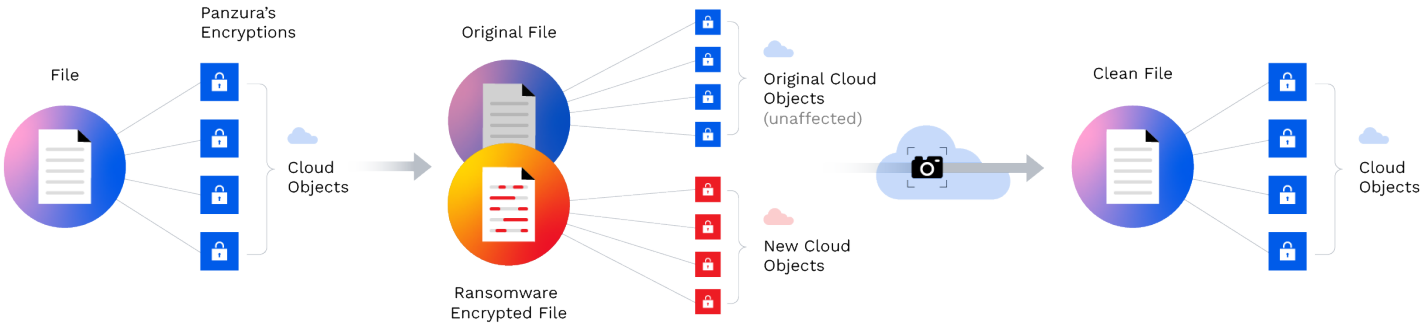

Our approach to immutable storage naturally allows rapid recovery of pristine file data from immutable snapshots taken at granular and configurable intervals. The speed at which CloudFS captures and moves file data into the centralized object store from anywhere in the world lets organizations achieve a recovery point objective (RPO) of under 60 seconds on a global scale. Individual files, folders, and even the entire file system can be rapidly recovered to a precise point in time immediately preceding data damage.

Moreover, users can recover saved versions of files through integration with Windows Previous Version. This effectively provides a “Control Z” that doesn’t require IT intervention and allows users to maintain productivity even in the face of accidental deletions.

As a result, CloudFS customers can recover from malware, misconfigurations, and malicious or accidental data deletion quickly and return to productivity with minimal disruption. This also has remarkable implications for data security, integrity, and management as the industry continues to evolve.

Business Continuity in an Interconnected Landscape

As dependence on cloud services continues to grow, IT leaders are increasingly aware that misconfigurations and outages out of their direct control can have a severe impact on their business operations.

In the context of cloud storage, an outage can make data unavailable to users and processes and can represent a single point of failure that does not meet the organization's requirements for resilience.

Panzura CloudFS is uniquely capable of bolstering storage resilience by allowing organizations to mirror their file data to a second object store, which may be in the cloud or on-premises. With cloud mirroring, new and changed file data is written to two object stores simultaneously, ensuring both have an identical copy of the data held in the file system.

During normal operations, the primary object store immediately serves up data to users and processes as required, while the secondary object store accepts new and changed data only. If the primary store becomes unavailable, file operations switch to the secondary object store, which remains in operation until normal service is restored.

Once both object stores are operational again, new and changed data is synced to the object store that suffered the outage so it becomes consistent once more. With cloud mirroring enabled, organizations effectively double their data availability and safeguard business continuity.

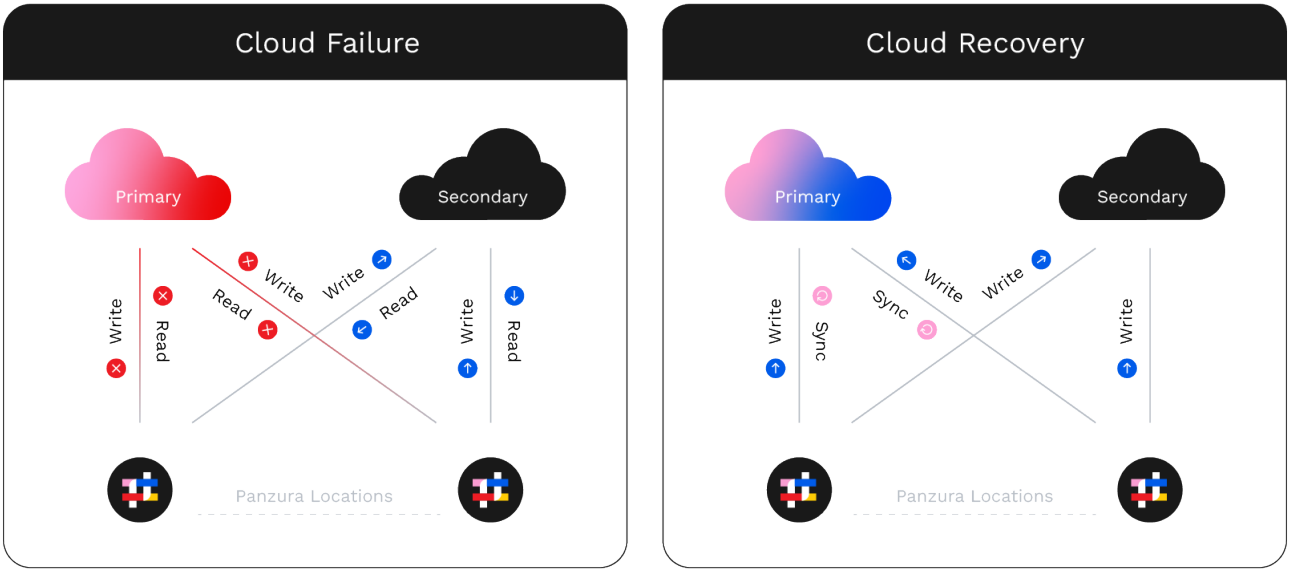

PART TWO

Command and Control: Mastering Datasphere Complexity to Accelerate Business Results

In today's fast-paced digital landscape, enterprise IT managers are confronted with the complex challenges of managing extensive storage volumes, escalating costs, and ensuring seamless data control. Exponential growth emphasizes the magnitude and potential impact of unstructured data on our future. More than 50% of organizations experience data sprawl and bloat. The average employee uses 37 tools for daily work, and 70% of those are used for unstructured data.

The ability to command and control data effectively is crucial for maintaining operational efficiency, ensuring data integrity, and optimizing storage resources. Panzura delivers file data management solutions that not only address these challenges but also deliver measurable business advantages.

Panzura's technology is built for the cloud, offering a hybrid cloud model that makes secure, global cloud file management accessible for enterprises. It's a radical rethink of how unstructured data is handled, how value can be extracted from it, and how the status quo of legacy file and object storage systems is no longer fit for purpose in an era of artificial intelligence (AI) and a borderless workforce.

Among the early and distinctive advantages of Panzura technology has been the ability to transform cloud storage into a global file system. This allows large matrix organizations to use the cloud as a high-performance, globally available data center. A centralized command-and-control structure allows organizations to deduplicate and consolidate data, achieving a reduction in overall unstructured data storage footprint of up to 80%.

In fact, Panzura data management platforms can together be considered a nascent data fabric. As an architecture and set of data services that provide consistent capabilities across a choice of endpoints spanning hybrid cloud environments, Panzura is on a path to simplify and integrate data storage, enabling seamless data access, sharing, and governance.

Data fabrics and real-time analytics consume internal, external, and edge data to meet future IT and business demands. However, 43% of IT decision makers fear their IT infrastructure won't handle future data requirements. Interoperability is at the heart of Panzura's innovation strategy, ensuring that data is always routed to the right storage solution and optimizing workflows.

Addressing data growth and cost challenges

Command and control ensure that enterprises can manage their data efficiently and securely. For instance, Skyscanner, a trailblazer in the travel industry, leveraged Panzura's global file system for secure and reliable file data management. They have achieved not only lower operational costs but also better productivity through immediate file consistency and local-feeling file performance for every user, regardless of location.

Other organizations such as a global Paytech provider achieved an average of 70% reduction in data volume across file data stored in both AWS and Microsoft Azure, through granular file deduplication and compression.

Panzura's comprehensive data services platform Symphony complements these command-and-control capabilities. It allows data to be managed and mobilized by any attribute, making complex data governance rules easier to apply. Organizations can take action and route assets to the right location at the right time, for the lowest possible cost.

This comes into sharp focus when considering that the sheer volume of data generated and stored by enterprises has been growing at an astounding rate. According to IDC, the global datasphere is expected to reach 175 zettabytes by 2025. This rapid growth poses significant headwinds for IT and operations teams tasked with managing storage infrastructure.

Effective command and control of data becomes increasingly critical in this context. With the increase in data volume, storage costs can quickly spiral. Traditional storage solutions often require considerable capital expenditure for hardware, software, and maintenance. Additionally, the ongoing operational costs associated with managing and scaling these systems can strain IT budgets.

Panzura addresses this issue by leveraging the power of cloud storage. By shifting data to the cloud, organizations can take advantage of scalable, cost-effective storage solutions that are widely recognized to reduce total cost of ownership (TCO) compared to on-premises storage. This shift not only provides cost savings but also enhances the command-and-control capabilities of IT managers, allowing them to handle data more effectively and maintain better oversight.

Fluor relies on Panzura file data technology in CloudFS to boost security, yield, and optimize storage. The firm faced challenges with high latency between different sites, which Panzura technology effectively addressed by providing real-time file locking and synchronization. This not only improved collaboration across multiple sites but also resulted in financial savings.

Enhancing data visibility and collaboration

Panzura's data mobilization and management capabilities also play a crucial role in addressing data growth and cost challenges. By providing complete interoperability between file system deployments and various storage environments, the data services platform helps teams control and inspect unstructured data across private storage and the cloud. Data is always routed to the most cost-effective storage option or tier, further enhancing command and control over vendor and hardware lock in.

Another critical challenge faced by IT teams and business leaders is gaining visibility over their data. Traditional storage systems often result in data silos, making it difficult to have a comprehensive view of the entire data landscape.

This lack of visibility can lead to inefficiencies, redundancy, and compliance risks. Moreover, up to 80% of enterprise data remains largely untouched or unanalyzed after its creation. Unstructured data accounts for a vast majority of data collected, yet only half of it is analyzed.

70%

reduction in total cost of ownership when moving from on-premises storage to cloud storage

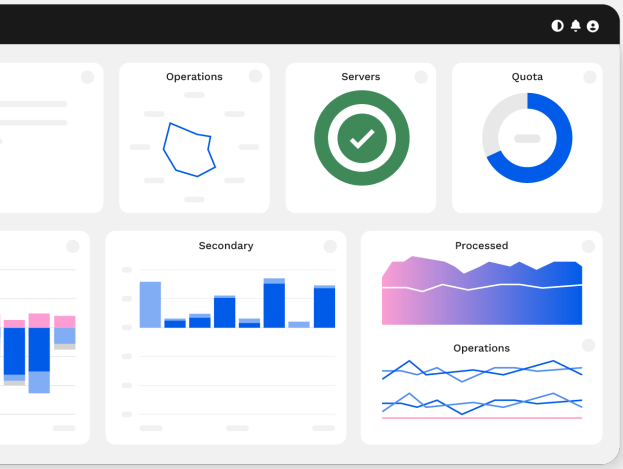

Panzura provides data management and services that offer real-time visibility, control, and data operations for all unstructured file and object data across the organization. These capabilities are designed as unified, “single pane of glass” solutions that allow IT managers and line-of-business (LoB) leaders to monitor, manage, and control data with ease.

With advanced analytics and reporting capabilities, organizations can gain valuable insights into data usage patterns, identify potential bottlenecks, and optimize storage resources. This enhanced visibility is a key component of effective command and control of data, enabling IT managers to make informed decisions and maintain data integrity.

The rise of AI in the data equation

As more and more data are processed through AI workflows and pipelines, including Large Language Models (LLMs), this is of critical importance. Nearly 80% of new data pipelines will be for ingesting, processing and storing unstructured data. Organizations that fail to adopt and integrate AI technologies into their data workflows risk falling behind, jeopardizing their market position and long-term viability.

Addressing these challenges, one of the standout features of Panzura is its global namespace capability. This allows organizations to create a single, unified view of all their data, regardless of where it is stored. A global namespace ensures that AI models and workflows can access data seamlessly. This eliminates the need for complex data migrations or multiple data access points, streamlining the integration process.

It also enables immediate file consistency across all locations. This is crucial for AI workflows that rely on real-time data processing and analysis. Consistent data ensures that AI models are trained and operate on the most up-to-date information, improving accuracy and reliability.

As AI workflows often require processing large volumes of data, the scalability offered by a global namespace allows organizations to efficiently manage data infrastructure. Similar to complex enterprise workstreams, AI projects inexorably involve collaboration across different teams and locations. A global namespace allows seamless data sharing and collaboration, ensuring that all stakeholders have access to the same data sets.

Ensuring data accessibility, therefore, must be balanced with command and control over data protection and security. An Illinois hospital exemplifies this balance. Their implementation of the Panzura hybrid cloud file platform thwarted a potentially catastrophic ransomware attack. Leveraging the inherent immutability of object storage, Panzura technology ensured critical patient data emerged untouched and accessible.

Panzura's data services platform takes this a step further. It enhances accessibility by providing granular control over unstructured data. ITOps can define, allocate, and optimize file and object workflows, so data remains safe and immediately available. This level of control is essential for maintaining data integrity and facilitating seamless collaboration across teams.

For instance, universities use Panzura technology to manage vast amounts of academic and research data. This has allowed them to maintain data integrity and make data readily available for academic collaboration and research.

Embracing file data management and services

Panzura's command and control capabilities also extend to regulatory compliance. By providing robust encryption and data protection measures and achieving FIPS 140-3 certification, Panzura ensures that sensitive information remains secure and compliant with industry mandates. This level of control is essential for organizations handling large volumes of confidential data, as it mitigates the risk of data breaches and unauthorized access.

Compliance with regulatory mandates around data is essential for countless organizations. Panzura's hybrid cloud data management solution not only ensures adherence to these stringent requirements but also fosters innovation and drives business results.

By transforming cloud storage into a more flexible, powerful business-centric asset, Panzura enables teams to work collaboratively across multiple sites with confidence that their unstructured data is always where and when it needs to be, ultimately enhancing productivity and operational efficiency.

As the enterprise continues to navigate a more competitive world, the benefits of command and control over data are clear. Panzura's commitment to data excellence and outstanding data experiences makes it the ideal command and control partner for organizations looking to connect their vantage point in the global datasphere to real and lasting business results.

PART THREE

Immediacy: Data Delivered on Demand When it's Needed, Where it's Needed

Immediate data delivery to users and processes without making file copies gives organizations a powerful competitive advantage. And yet, it's a challenge that very few solutions come close to addressing.

A move to cloud storage is often driven by a desire to consolidate multiple localized Windows File Shares and achieve greater resilience for data. What is not immediately obvious is that some of the functionality inherent in Windows File Shares is tightly tied to network-attached storage (NAS) and as a result, does not work in the same way when storage is moved to the cloud.

Organizations accustomed to working from local file storage generally enjoy a high performance file environment. There, file data is close enough to users and workloads that file operations can be as performant as the software application, file size, and complexity of any file dependencies will allow. In this scenario, a simple Microsoft Word document will respond almost instantly, whereas a large and complex design file is likely to take a few moments to open and save.

However, for organizations with multiple locations, timely data delivery to people and processes has always posed a significant challenge. The ideal scenario from the perspective of cost, productivity, and efficiency is to consolidate file storage and allow the entire organization to work from a single data set that can be appropriately protected and maintained.

Achieving this is one of the hardest jobs IT leaders face. Distance and latency go hand in hand and as users, applications and data become ever more widely distributed, data delivery becomes exponentially more difficult.

Poor file performance is an insidious time suck that steals seconds and minutes of productivity from every user at every file interaction. These can seem trivial on an individual basis and this lost productivity can go largely unnoticed unless users make a point of repeatedly bringing it to the attention of the IT team.

When multiplied across the organization over the course of a year, these lost seconds and minutes represent a painful reality for organizations vying for competitive positioning. Quite simply, they are the price organizations pay for poor data delivery. The costs can be staggering.

As soon as file operations hamper the ability of users to produce their best work, the organization has a problem. Now, employees have stopped operating at the speed of thought and have begun to battle frustration at best, and lost work at worst.

For the business, it's a triple whammy of costs. There's a base salary paid with the expectation of a certain number of productive hours worked, while the available tools are standing in the way of productivity. There's the cost of those tools, which are behaving in a way that is not in the best interests of the business.

The cost of lost opportunity is the greatest cost of all because it extends far beyond the seconds, minutes and hours directly lost due to slow file operations. Instead, the frustration and distraction it causes have a direct impact on the quality of work employees are capable of both in the moment and long term.

The resulting insidious decline in performance can negatively impact productivity and quality of output on an ongoing basis. It means the organization may never fully realize the value their employees have the potential to offer, individually or collectively.

The same scenario holds true at every level of the organization. In the digital age, almost every organization has a substantial dependency on data, which means poor file performance is likely to negatively affect most if not all areas of the business. It's a problem that's imperative to solve.

Organizations that opt for substandard file environments fall behind their competitors because they cannot consistently position their best people to do their best work. Worse, high performers who are driven to do good work actively seek an environment in which they can thrive, so organizations risk losing top talent to their competitors.

Only immediate is fast enough

User expectations of digital experience have never been higher. In our personal lives, we've become accustomed to ordering entertainment, meals, products, and services for immediate delivery. When we book an Uber, we expect our ride to arrive within minutes — we don't expect and won't tolerate the application taking minutes to open and respond. In fact, when digital experiences fall short of our expectations, we typically respond by using a competitor's service.

At work, these same expectations hold true. Savvy IT leaders understand this paradigm shift means that users have a fraction of the tolerance for poor file performance that they had just a short time ago. Now, they become frustrated with their file experience very quickly and this frustration affects their work.

This brings us to a critical collision-avoidance feature of Windows File Shares that is frequently overlooked as organizations migrate to the cloud.

The role of file locking is to avoid data corruption and conflicts by preventing multiple users from editing the same file simultaneously. Without it, users can overwrite each other's work with no indication that they are doing so. The set of edits that is saved last will win. Other edits will be lost. Hours of expensive work can be undone with little to no ability to recover it, even with IT intervention.

For this reason, file locking capabilities date back to the earliest multi-user operating systems. Over the years, they have become progressively more granular and this has empowered significantly greater user productivity.

The bluntest mechanisms for locking restrict read-write access to one user at a time, while allowing other users to open the file as read-only. The most precise capabilities enable co-authoring, allowing multiple users to edit a file at the same time without overwriting each other by locking the portion of the file they are working on. When Google first launched Google Docs, Sheets and Slides, it emphasized their collaborative features and real-time editing capabilities as key ways to improve team productivity.

Microsoft also offers granular range-locking capabilities in Excel, Word, and PowerPoint, again drawing attention to productivity gains that are enabled by real-time collaboration. Co-authoring is available with Microsoft 365 subscriptions. Developers of other software applications have followed suit and as a result, file collision avoidance mechanisms have become ubiquitous — they simply work. Until of course, organizations move to the cloud and put significant distance between file storage and the users who work with those files.

Now, latency stands in the way of real-time collaboration. It takes much longer to open and save files than it previously did. Response time is so slow that complete file locking mechanisms can fail and range locking has no hope of working, so productivity stalls.

Myriad solutions attempt to address this, including those developed by cloud storage providers in order to allow organizations to work with cloud-stored files. However, the list of solutions that come close to empowering a truly local-feeling file experience is short.

Panzura's hybrid cloud file services platform CloudFS stands alone in its ability to empower distributed organizations to use cloud or other object storage as a high performance, globally available data center.

CloudFS is underpinned by the fastest global file system on the planet, uniquely capable of immediate global data consistency regardless of the number of locations or users. That means that users and processes can access and work with up-to-the-minute file data regardless of when it was last changed, and the distance between them and the user or process that last changed it.

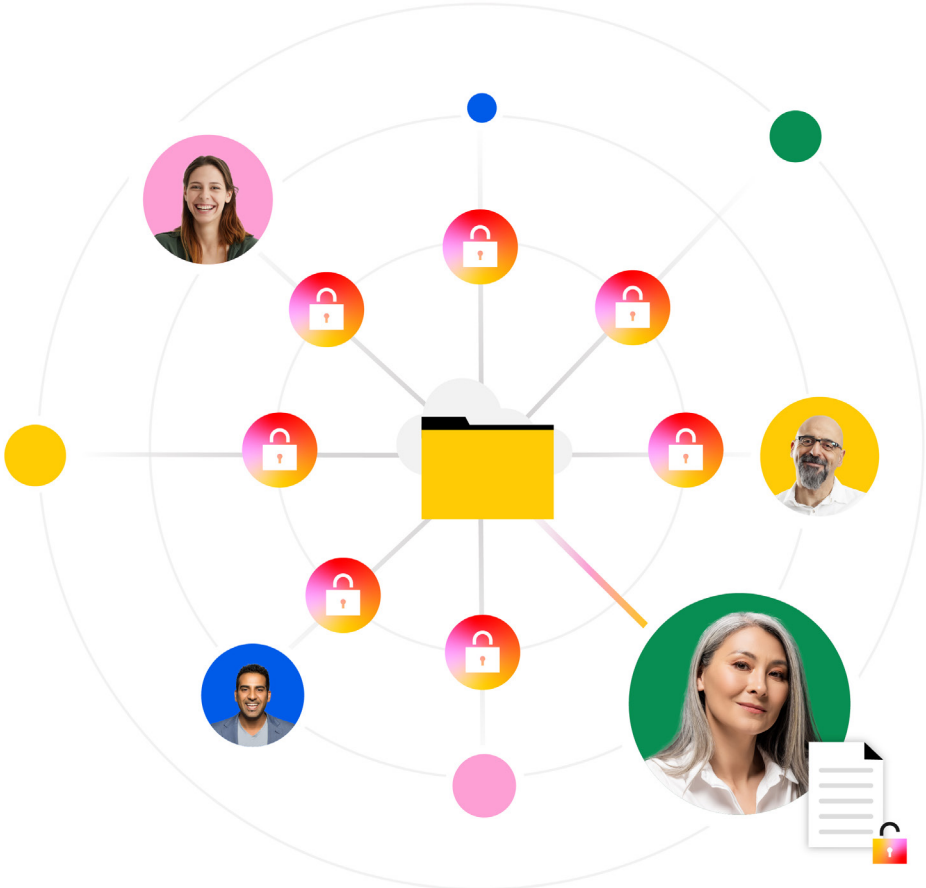

For example, global engineering firm AFRY has empowered their 16,000-strong workforce to co-author some of the world's most complex CAD design files in real-time, creating a high-performance file environment that allows them to work as if everyone is in the same office. In reality, AFRY spans 4 continents.

In addition to the exception level of control this gives AFRY's ITOps teams over their expanding data sets through patented, granular deduplication and compression, this approach to immediate data delivery creates a competitive advantage that is very hard to beat.

> CloudFS makes data resilient to damage, whether due to accidental editing, deletion, ransomware, or malware. Read-only snapshots provide point-in-time recovery for individual files and folders, or the entire file system if required.
>
> **BANKING & FINANCIAL INSTITUTION**

As organizations grow, they often sacrifice agility as their size forces them to move and respond more slowly. With Panzura, even global organizations can be as nimble as their smaller, local competitors, with a file environment that supports and fosters real-time collaboration without the risk of file collisions.

While Panzura is well known for enabling immediate file delivery and real-time collaboration with some of the most complex applications and largest files, organizations within all industries realize extraordinary benefits from immediate file data delivery. The fact is, organizations that makes users or processes wait to access files on a consistent basis struggle with productivity and this stands in the way of growth and innovation.

Further, inability to feed AI with curated data on demand is likely to widen competitive gaps as organizations that can make their data work for them by putting it where it needs to be, when it needs to be there pull away from those that cannot.

In a data-driven world, the winners are those who master data management.

Digital transformation and IT infrastructure modernization projects seek to position organizations to leverage the best of cloud technology and yet, IT and line of business leaders often find themselves overwhelmed and confused by the myriad solutions available to them. It is increasingly clear that the benchmark against which these solutions must be measured should consider three key aspects above all others.

The solution must give the organization data resiliency, so that if and when the organization's defenses are penetrated and data is damaged, pristine data is available to be restored promptly without loss or unnecessary delay.

Secondly, the solution must provide command and control over spiraling data volumes (and therefore associated costs). This must encompass a complete view over all data, regardless of where it is generated or stored, and it must empower IT leaders to take action on the insights that are surfaced, within the same control plane.

Equally, the solution's approach to file services must control the rapid growth of unstructured data through granular deduplication techniques that strip out duplicated data blocks. Solutions capable of clever use of metadata can be used to "point" to the data blocks that comprise a file instead of allowing files that are substantially similar (and often almost identical) to occupy storage space with data that is redundant because it has already been stored.

Finally, only immediate data delivery is fast enough to empower organizations to achieve maximum productivity. The cumulative impact of wasted seconds and minutes on the organization's competitive posture cannot be understated.

Organizations that achieve resilience, command and control, and immediate file delivery across distance, regardless of their industry, gain competitive advantage that is very difficult for others to surpass.