

White Paper

Panzura: Furthering the Potential of Public Cloud Services

A High Performance Global File System Built on a Hybrid Cloud Storage Infrastructure

By Scott Sinclair, ESG Senior Analyst

August 2016

This ESG White Paper was commissioned by Panzura and is distributed under license from ESG.



Contents

Introduction	3
The Power of Cloud Services.....	3
When The Cloud Becomes Yet Another Copy.....	5
The Need for Efficient Hybrid Cloud Storage: Hot Edge and Cold Core.....	6
Introducing Panzura’s Hybrid Cloud Storage	7
The Bigger Truth.....	8

Introduction

The rise of cloud services is not a temporary phenomenon. It is here to stay, and is fundamentally changing the way organizations procure, deploy, and access information technology. And the benefits that organizations receive can be dramatic. Though the cost-effectiveness of cloud infrastructure is commonly understood as the biggest benefit, there are other, potentially more valuable benefits available, such as increased agility.

There is a reason that economists and accountants speak of money in terms of present value—there is a time component to value and \$100 today is worth more than \$100 two years from now. Similarly, a workload that can be deployed in an hour is worth more than that same workload deployed two months from now. However, on-premises infrastructure takes time to architect, procure, install, and provision. In other words, the time it takes to deploy infrastructure on site is starting to cost businesses too much.

Much of the challenge with on-premises infrastructure design stems from the perpetual onslaught of data growth. Faced with rising capacity levels and an ever-increasing need to support new workloads, onsite IT administration teams are hard-pressed to keep up. Seeking a solution, IT organizations are turning to public cloud services. Yet, integrating the public cloud adds its own level of complexity. The challenge of the physical separation between on- and off-premises becomes front and center. The latency of the WAN may be the most significant hurdle organizations face when adopting public cloud services. Moving data to and from public cloud services takes longer and each transaction increases the cost of using the cloud. This challenge of physical isolation can often give IT organizations pause when it comes to executing a cloud strategy, or limit which data sets can leverage public cloud services.

Fortunately, hybrid cloud solutions exist to mitigate these challenges. Combining the benefits of on- and off-premises technology, these solutions can ease the integration of public cloud resources while optimizing data movement over the WAN. Panzura, a leader in hybrid cloud storage solutions, offers a number of innovations specifically designed to maximize the potential of hybrid cloud architecture to serve global or multi-site businesses.

The Power of Cloud Services

The integration of public cloud services currently resides at the heart of many IT data center strategies. In a recent research study investigating IT spending intentions, ESG polled 633 IT decision makers. When these IT leaders were asked to identify which measures their organizations would take to reduce or otherwise contain IT expenditures over the next 12 months, they most commonly identified the increased use of cloud computing as an alternative to in-house infrastructure. Its rank as the most-cited response in 2016 represents a seven-year journey up from the number nine overall spot in 2009.¹ The key takeaway is that, over the past seven years, IT leaders have learned the cost-effective potential of public cloud services.

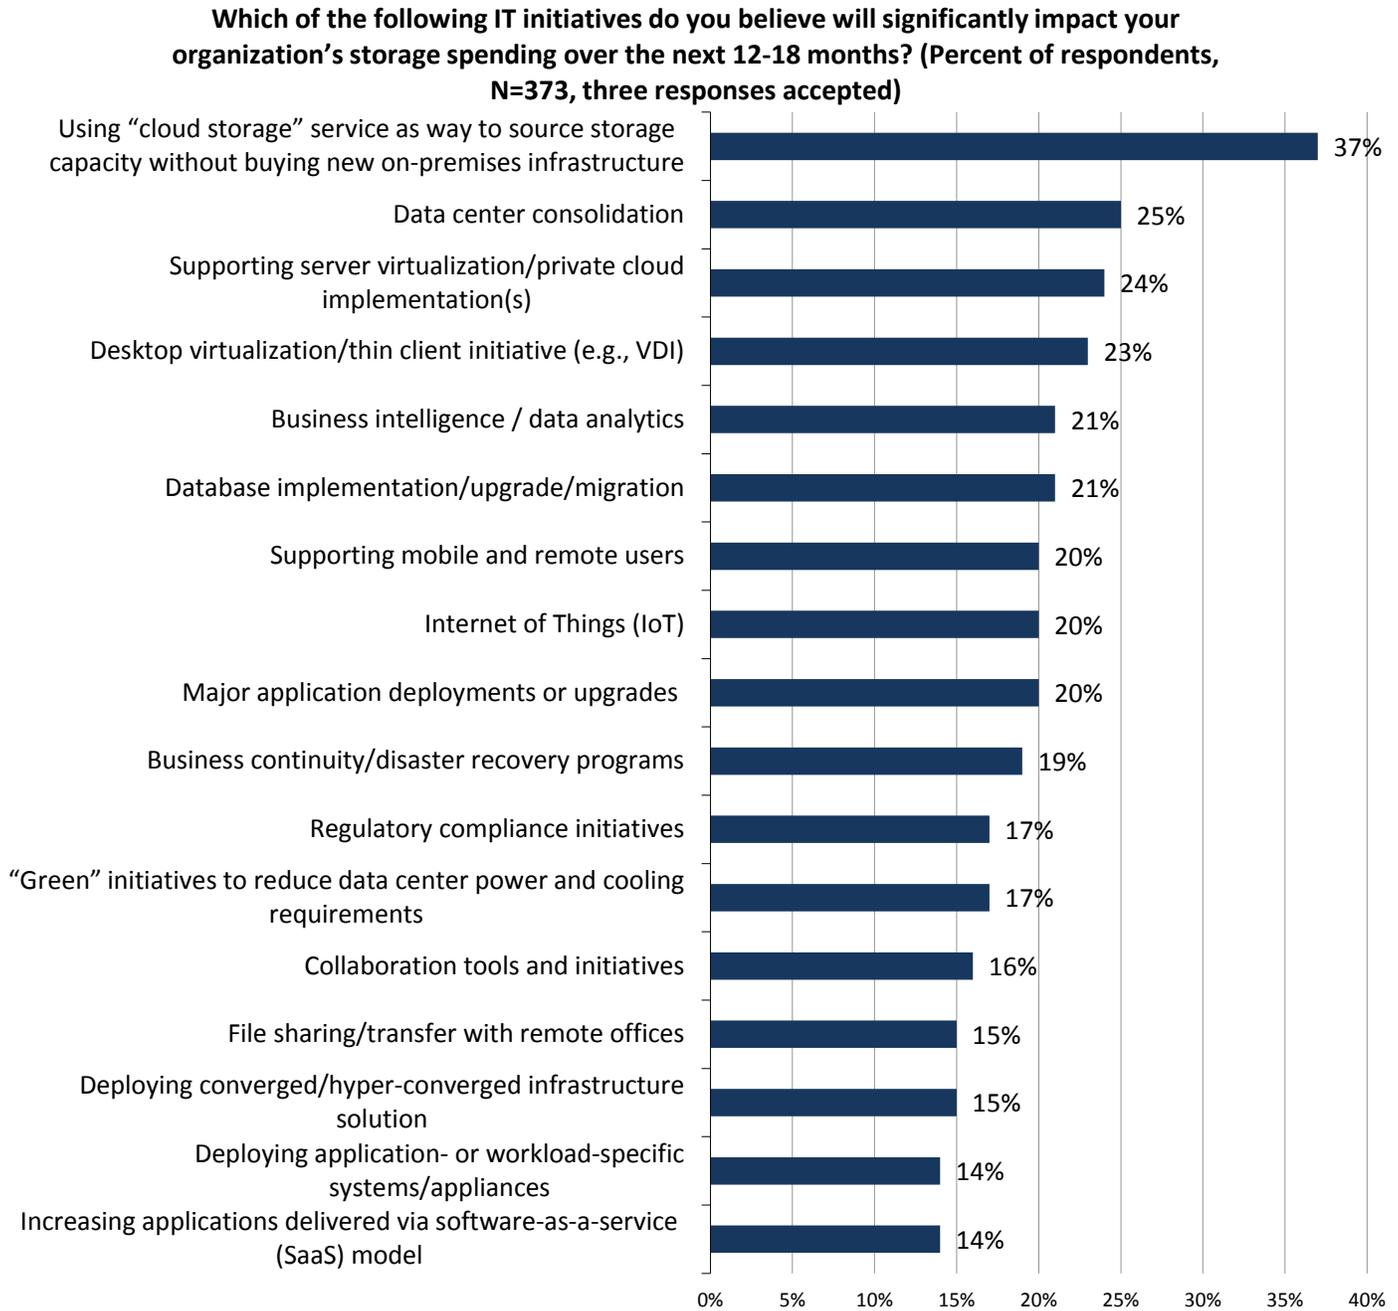
Storage administrators are also seeing the potential. As part of a separate investigation into general storage industry trends, ESG surveyed 373 storage decision makers responsible for their organizations' data storage infrastructures. This research study investigated multiple storage technologies, including the public cloud. When storage administrators were asked to identify the initiatives expected to most impact their storage spending over the next 12 months, the most commonly identified answer was the use of cloud storage services as a way to source storage capacity without buying new on-premises infrastructure. This response was dominant, chosen by 37% of respondents—a response 12 points higher than the second most commonly identified answer (see Figure 1).² This 12-point delta is greater than the separation between

¹ Source: ESG Research Report, [2016 IT Spending Intentions Survey](#), February 2016.

² Source: ESG Research Report, [2015 Data Storage Market Trends](#), October 2015.

the second and the 17th most commonly identified answer. In other words, regardless of the IT audience, or the research study, cloud is making an impact.

FIGURE 1. Business Initiatives with the Greatest Impact on Storage Spending in 2016



Source: Enterprise Strategy Group, 2016

However, despite the benefits, challenges do exist, often stemming from the physical separation of onsite and offsite resources. There is a limit to the speed at which digital bits can be transported from one location to another. In an ideal setting, IT infrastructure would be designed so that the data always resides as close to the application as possible in order to minimize the latency incurred. This is expensive, however, and budgets are finite. In this era of exponential data growth, placing data next to the application is rarely feasible for every workload. The increased latency introduced with public cloud services slows each transaction, reducing the effectiveness of local applications. To further increase costs, many

public cloud services charge for data transactions, so each data access across the WAN not only takes longer but also drives the cost of storage higher. In addition, this latency challenge impacts the ability to move data.

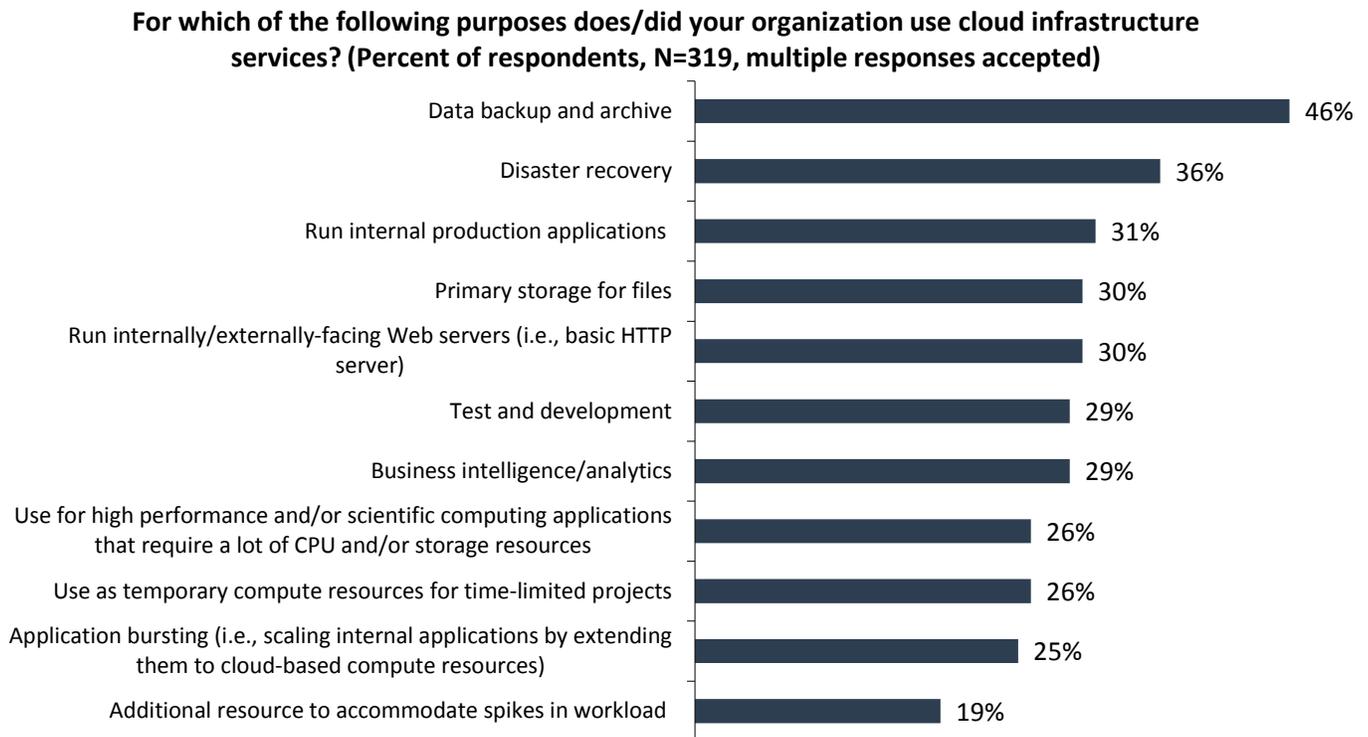
Migrating a data set to the cloud can be time-consuming and costly, often adding a level of permanence to on- versus off-premises design decisions. Ultimately, the challenge of distance limits the agility offered by public cloud services. It may take only moments to stand up an application in the cloud, but moving the requisite data may take days or even weeks. Once the data has been moved, organizations often find that the on-premises content has changed during the time it took to move data to the cloud. This adds an additional effort to reconcile the data set, increasing costs.

When The Cloud Becomes Yet Another Copy

In response to the limits created by the latency challenge of the WAN, IT organizations often take a measured approach to cloud adoption resources, with a common use case being to store yet another copy of data. In ESG’s 2016 IT spending intentions research, IT leaders whose organizations were leveraging public cloud services were asked to identify the purposes behind their use of the cloud. The most commonly identified response was that organizations use cloud infrastructure services for data backup and archive, followed by disaster recovery (see Figure 2).³ In other words, an often-popular use of the cloud is to serve as an incremental and redundant copy.

While somewhat understandable in a risk-averse industry such as IT, leveraging the cloud solely for the purpose of retaining an additional copy limits the benefits of cloud services and increases data center complexity. The net result serves to add another storage silo, albeit an off-premises one, to manage, maintain, and protect. Additionally, this new silo is separated from the data center by the latency of the WAN and the increased cost of data movement.

FIGURE 2. Purposes for Use of Cloud Infrastructure Services



Source: Enterprise Strategy Group, 2016

³ Source: ESG Research Report, [2016 IT Spending Intentions Survey](#), February 2016.

The Need for Efficient Hybrid Cloud Storage: Hot Edge and Cold Core

Embracing public cloud services to maximize the benefits of cloud while minimizing the cost and complexity requires an efficient hybrid cloud storage implementation. A leader in hybrid cloud technology, Panzura offers businesses a new paradigm to adopting public cloud services, referred to as a hot edge and a cold core. This model abandons the antiquated concept of the cloud serving as another storage silo, one even more isolated than those that exist on-premises. The core concept of this architecture is that the public cloud serves as the central repository for all data storage. The public cloud infrastructure serves as the “cold” core of the storage ecosystem, where “cold” refers to the storage of less frequently accessed data sets. The “hot” edge, that serves the “hot” or active data, is delivered via high performance local data caching appliances placed at each location. This architecture, when combined with global file locking and global deduplication, is designed to drastically limit the impact WAN plays in the ability to utilize public cloud resources. As a result, leveraging the public cloud as the core data repository, while maintaining active data caches on site, delivers a number of benefits, including:

- Significantly reducing the amount of storage infrastructure required to manage, support, power, and cool at each site.
- Minimizing the amount of unnecessary copies of data distributed across sites.
- Minimizing the amount of data that traverses the WAN.
- Offering workloads at each site and in the public cloud access to the same consistent data set, enabling workload agility and data mobility.

Leveraging this hot edge and cold core architecture, Panzura’s hybrid cloud technology is then able to deliver a number of business benefits, for example:

- **Simplify Public Cloud Adoption:** Panzura’s technology enables cloud storage to present the look and feel of on-premises enterprise storage. By running a Panzura controller in the cloud, local applications can scale data into public cloud services without complex configuration. The result removes the complexity of managing multiple on- and off-premises storage silos.
- **Alleviate the Pain of Data Growth:** In ESG’s research into general storage industry trends, storage administrators were asked to identify the biggest challenges with regard to their storage environments. The ability to scale data storage into the public cloud directly addresses seven of the top ten challenges identified by respondents in the research study. These challenges include hardware costs (27%), the rapid rate of data growth (26%), data protection (26%), staff costs (23%), running out of physical space (20%), power and cooling costs (18%), and device management (18%).⁴ Ultimately, procuring, deploying, managing, and servicing incremental storage frames in an effort to keep pace with data growth is a complex and costly battle that Panzura’s hybrid cloud technology can help resolve.
- **Consolidate Workloads Via Intelligence and Automation:** Performance demands often determine the cost of storage. Higher performance often demands a higher cost per capacity. Yet, the performance demands of specific workloads and data sets can change over time. Managing storage tiers and ensuring that the right data set is located on the optimal storage resources is a challenge even for on-premises storage environments. Managing

⁴ Source: ESG Research Report, [2015 Data Storage Market Trends](#), October 2015.

storage tiers across on- and off-premises resources can be even more complex. Panzura's solutions present a common storage interface yet automatically manage performance with intelligent caching, further simplifying storage architecture design.

- **Provide for Workload Agility:** With a consistent data set being maintained on the public cloud, all of the independent sites, as well as the cloud, are able to access the same consistent data. Workloads can be deployed at a certain site, or in the cloud, based on what makes sense for the business, rather than the location of the data.

The value of a hybrid cloud solution is achieved through multiple aspects: deliver the simplicity of local storage, enable the agility of public cloud services, and empower the solution to serve a wide breadth of workloads. Panzura's hybrid cloud storage technology is able to follow through on each of these promises with innovations that deliver simplicity through automation and reduce costs through efficiency in data storage and movement.

Introducing Panzura's Hybrid Cloud Storage

Panzura's hybrid cloud storage technology encompasses a number of technological elements designed to dramatically simplify storage resource management while simultaneously optimizing data placement and movement for a more efficient ecosystem. The result is a hybrid cloud storage solution that significantly reduces the impact of WAN latency. Panzura's technology offers several technical innovations critical to delivering these benefits, including:

- **Globally Consistent and High Performance Metadata:** Metadata is data about data. Fast metadata access is critical, since a majority of file system interactions are often metadata actions. Searching, traversing the file tree, and identifying changed files all require metadata access. The actual opening of a file and modification of the data occur far less frequently. Additionally, for multi-site organizations, metadata must be consistent in order to ensure that files are accessed safely and securely. Panzura's technology stores metadata separately from the data, locally in flash, and all file metadata operations run locally to the user. Panzura then maintains consistent metadata synchronized across all controllers. File locks are distributed across all of the nodes and locking metadata is updated immediately across all controllers. By separating the file locks from the files, Panzura places the lock closer to the user, eliminating the effect of WAN latency. This technology delivers an immediately consistent file system across locations with no separate versions to manage, like some other solutions in the industry.
- **Global Inline Deduplication:** Panzura offers inline deduplication to minimize the amount of data that traverses the WAN. Each onsite controller has the full deduplication table of all files in all controllers. The net result is a massive scale, multi-site deduplication that optimizes data capacity and movement. Less data traverses the WAN, reducing costs and speeding up transfers. Once data is in the cloud, it requires less space and less cost. The technology also supports deduplication-aware snapshots for further capacity savings.
- **Policy-driven Caching and Fast Sync:** Panzura's policy-based caching technology is configurable and allows the same system to serve both transactional and archive workloads. In addition, Panzura combines machine learning based on file usage with the ability to "pin" data to determine the data that should reside in the on-premises

Feature Overview

- **Global File System with global file locking**
- **Global Inline Deduplication**
- **Integrated Global Snapshots, Backup, and Disaster Recovery**
- **Policy-driven Caching and Fast Sync**
- **Data Integrity with End-to-end FIPS 140-2 Security**
- **Ability to run apps locally or in the cloud against the same data sets**

cache. Local caching of the file helps reduce the amount of data that needs to be transferred, but often the time to transfer the data is significantly less than the time for the thousands of file operations to traverse the WAN. For some applications, simply opening a file can generate thousands of file operations before the application opens. Without global file locking, each of those operations would have to traverse the WAN and back, multiplying the latency impact. Within global file locking, Panzura is able to lock files at the byte range level. This allows for multiple users to access the same file simultaneously from globally separate sites, eliminating the need to generate or distribute multiple versions. Without the ability to lock content at the byte range level, applications that open for write would immediately fail when trying to access a file in use by another individual.

- **A NAS Mount to Directly Access Data in the Cloud:** To maximize the agility offered by the hybrid cloud, businesses must be free to deploy new workloads on- or off-premises. Panzura's technology provides consistent data access across multiple nodes located across the globe. Additionally, by offering a NAS mount in the cloud, companies are able to deploy and run applications in the cloud without requiring an additional data migration step or needing to alter the application to communicate via a separate cloud protocol. This allows applications to be deployed at any one of multiple sites or off-premises. If offices are less than 10ms away from the cloud data center, companies can eliminate all unstructured data from those offices and use the in-cloud NAS for on-premises file access.

These innovations enable Panzura to serve a wide variety of use cases, such as global software build distribution, hybrid cloud NAS, high value asset distribution, and cross-site collaboration of technical applications such as CAD or building information modeling (BIM). Global multi-site organizations can deploy a globally consistent, nearly infinitely scalable file storage system. Additionally, by leveraging the cloud as the central site for data storage, global collaboration can be achieved by only moving the changed content across the WAN, increasing performance and reducing cost. Application development teams can stand up new workloads in the cloud or on-premises. The net result delivers a compelling level of increased agility, simplicity, and TCO savings for those organizations seeking to integrate public cloud services into their IT environments.

The Bigger Truth

Public cloud services possess the potential to deliver new levels of efficiency and agility to IT organizations. These benefits offer enticing opportunities to escape the firefighting and chaotic nature of on-premises infrastructure management. Yet, separated from the local data center, the value that the public cloud offers is limited. The chasm created by WAN latency limits the movement of data and the agility of workloads, and forces cloud migration activities to essentially become permanent ones. These limits take the potential of public cloud and effectively wall it off from the business.

Panzura's technology and the hot edge and cold core architecture design help eliminate the barrier created by WAN latency. Panzura's intelligent high performance caching appliances provide high-speed performance for off-premises content while automating the management of performance for individual workloads. Inline deduplication minimizes the cost of data transmitted and stored, while speeding up movement. Global file locking allows multiple sites around the world to work on the same content simultaneously without requiring multiple versions or copies. New workloads can be deployed on any site or in the cloud, providing them the ability to access the right content without having to migrate the data. These are just a few of the hybrid cloud capabilities made possible through Panzura's technology. Resolving the hurdles of the WAN, with technology like that from Panzura, is a key step in maximizing the benefit from the public cloud.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.



Enterprise Strategy Group is an IT analyst, research, validation, and strategy firm that provides actionable insight and intelligence to the global IT community.

© 2016 by The Enterprise Strategy Group, Inc. All Rights Reserved.

