

WHITEPAPER SERIES

**A NEW APPROACH IN ENTERPRISE
STORAGE SYSTEM DESIGN**

STRATISTOR™: UNLOCKING THE GLOBAL DATASPHERE



CONTENTS

1. EXECUTIVE SUMMARY

2. WHAT IS STRATISTOR™?

3. THE STRATISTOR ARCHITECTURE

4. CONCLUSION

EXECUTIVE SUMMARY

THE CURRENT (AND FUTURE) PROBLEM

Today's information technology landscape faces a monumental problem. Managing the massive amounts of data generated on a daily basis (i.e. online production data) across the world. As significant as it is, it is nowhere near the amount of data which has accumulated before today in the form of offline data (i.e. backups, archives). When you consider the total sum of all data globally, it is easy to conclude that current storage technologies are ill-equipped to manage this perpetually growing problem.

The data acquisition rate is at the core of the issue. In 2021, the global datasphere grew to over eight zettabytes (or eight trillion gigabytes) and is expected to grow to over well over 20 zettabytes by 2025.

In 2013, Eric Schmidt, former CEO of Google said,

“from the very beginning of humanity to the year 2003, an estimated 5 exabytes of information was created, which corresponds to 0.5% of a zettabyte. In 2013, that amount of information (5 exabytes) took only two days to create, and that pace is continuously growing.”

OPERATIONAL AND FINANCIAL

In modern enterprise environments, data is stored on physical devices such as storage arrays or SANs, which reside in the data center. Every day, more and more data is deposited on these devices. Eventually the inevitable hardware refresh is needed and all data must be migrated off to another newer device. This perpetually occurs throughout the lifecycle of the data set.

EXECUTIVE SUMMARY

Additionally, between refresh cycles, there are always components which need to be replaced due to failure or added to address new demands for capacity. When capacity demands exceed what the current platform can withstand, more hardware must be purchased and the data migrated. This cycle becomes a tremendous burden both operationally and financially for the organization.

Storage manufactures have recently attempted to shift the financial model from capex to opex. However, this proved difficult because creative ways were needed to address their long term hard costs. Even within these new opex models, there is still a significant amount of capex spend and often a complicated consumption component attached which ends up eroding the opex model benefits.

DATA CONTAINMENT

We cannot begin addressing the problem of scale without discussing the concept of parallelism. Historically parallelism was focused on delivering performance within a data storage system. However, we must now expand that to include parallelism in terms of the number of workers responsible for managing the dataset. In other words, we must divide the larger problem into smaller problems and concern ourselves both with performance and the addressability of the total data set.

This is precisely what StratiSTOR is intended to solve.

EXECUTIVE SUMMARY

WHAT'S CHANGED AND WHY NOW?

Some very interesting intersections have occurred in the global service provider world which has given rise to technology such as StratiSTOR. The first is the virtually unlimited network and storage capacity of the public cloud providers around the world. The second is the increasing availability of bandwidth and the dramatic drop in overall cost.

Many organizations today have made the move to 10Gb internet circuits with monthly costs now dropping below \$1,000 per month in enterprise datacenters.

According to industry analysts:

The global cloud storage market size is projected to grow from \$50.1 billion in 2020 to \$137.3 billion by 2025.

From 2008 to 2018, the cost per Mbps has decreased from \$9.01 to \$0.76. That is a 92% decrease over that time period and this trend continues today.

These conditions, combined with the right software, provide new and exciting ways to leverage public cloud storage providers not previously considered possible.

The basis for this whitepaper is to outline the high-level technology and economics associated with providing enterprise class storage using cloud backed storage endpoints from the various global cloud providers available today.

WHAT IS STRATISTOR?

StratiSTOR is a fully managed, cloud backed, software-based storage cluster which delivers enterprise class storage capacity and performance with cloud economics. It replaces the need to deploy traditional physical storage resources in the data center, avoiding all of the challenges and complexities associated with these types of traditional technologies.

StratiSTOR is designed to manage today's petabyte scale environments. It achieves this by leveraging the virtually unlimited capacity of the public cloud providers along with sophisticated parallelization technology to meet the demands and scale of today's modern workloads.

StratiSTOR is typically deployed as a series of virtual machines forming a high-performance highly available storage cluster, but can also be deployed as hardware with only the minimum requirements necessary to run the StratiSTOR operating system. In either deployment model, no mass storage devices of any kind are required as all storage capacity is cloud-backed from the global cloud storage providers.

WHAT IS STRATISTOR?

The features of StratiSTOR are designed to meet or exceed that of today's enterprise class storage platforms:

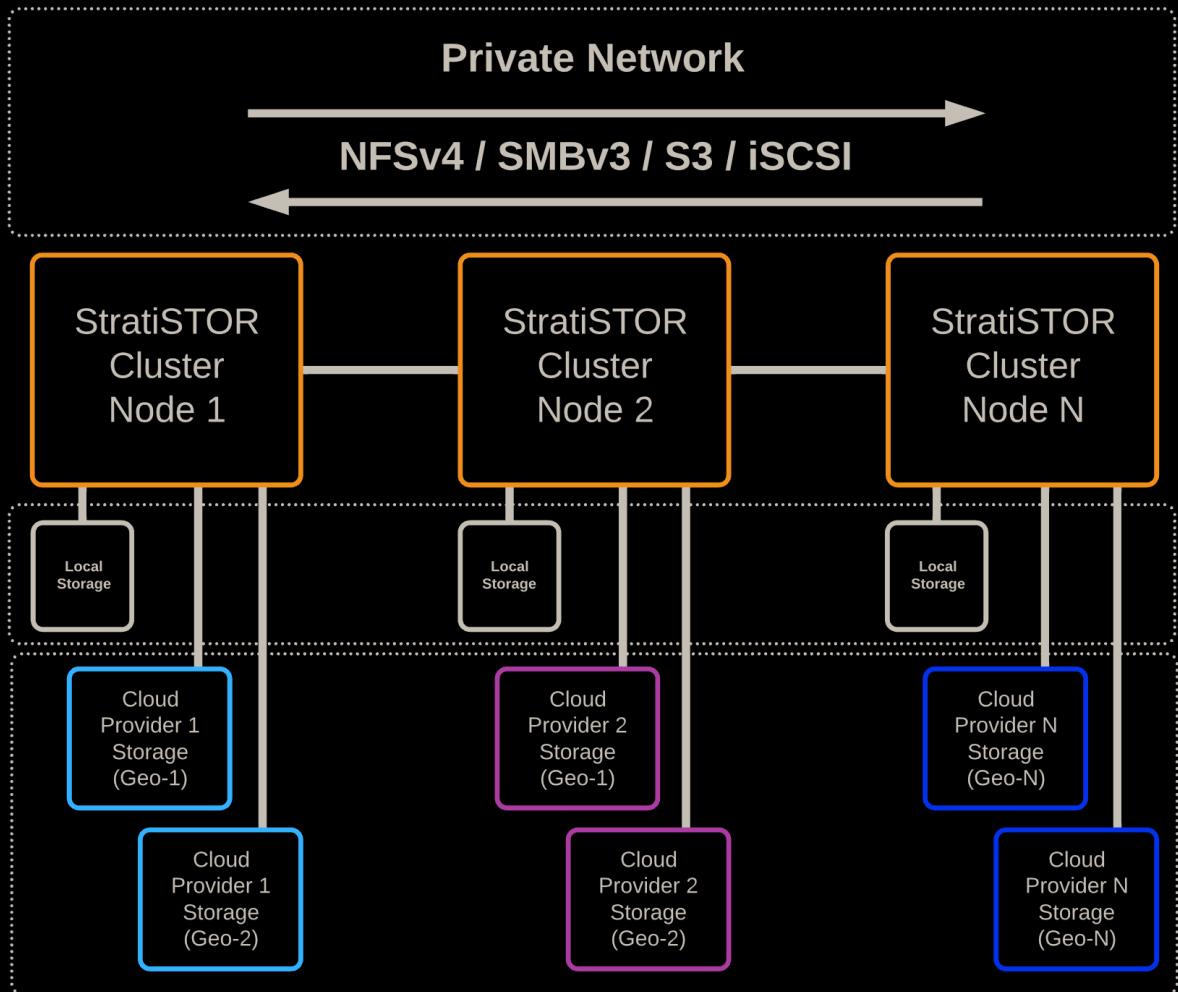
- Linear performance scalability as the number of nodes increase
- Non-disruptive scaling of nodes and/or capacity
- N+X high-availability
- Native in-flight and at-rest data encryption and integrity features to ensure your data is secure and error-free throughout its lifecycle
- Seamless integration leveraging universally compatible network protocols such as NFS, SMB, S3 and iSCSI

The benefits of StratiSTOR are extremely compelling:

- Fully managed and provisioned
- No mass storage needed
- Rapid provisioning and deployment
- Provider and geographically diverse
- Static monthly opex cost model
- No refresh cycles
- No migration of data needed ever again
- High performance (> GB's /sec)
- Compatible with all operating systems and applications

THE STRATISTOR ARCHITECTURE

StratiSTOR is deployed as a highly-available high-performance cluster across any arrangement of global storage cloud providers in any geography. The cluster is capable of scaling to hundreds of nodes as necessary to fulfill the performance and capacity demands within the environment.



PRIMARY USE CASES

StratiSTOR can be utilized in many different scenarios ranging from:

<u>TIER</u>	<u>STORAGE TYPE</u>	<u>APPLICATION TYPE</u>
Tier-1	production storage	databases, analytics
Tier-2	hot storage	file/web servers
Tier-3	warm storage	backups
Tier 4	cold storage	archival

Tier-1: StratiSTOR is specifically designed to provide the highest level of performance and availability for the most demanding applications, at scale.

Tier-2: StratiSTOR is powerful enough to manage the demand of enterprise concurrency achieving maximum workload density for enterprise environments.

Tier-3: StratiSTOR can scale to any level of capacity you need today or in the future for capacity hungry applications such as backup and business continuity applications.

Tier-4: StratiSTOR can manage large capacity and large retention applications such as archival data sets with ease.

CONCLUSION

The global demand for storage is accelerating to unprecedented levels and will continue to accelerate as our world becomes more connected. As the human pursuit for knowledge continues and our ability to explore expands, the amount of online data will grow geometrically. The scale of historical and archival data will continue to grow just as fast, to levels well beyond what traditional storage systems are able to address.

The scale of the global datasphere is such that there is no other way to address the problem without breaking it down into smaller more manageable components. StratiSTOR solves this problem by leveraging the cloud for what it was designed for with very compelling lifecycle economics.