FUTURUM.



IS THIS THE END OF TIERED STORAGE?

Q2 2022

DANIEL NEWMAN
Principal Analyst + Founding Partner

SHELLY KRAMER
Principal Analyst + Founding Partner

FRED MCCLIMANS
Senior Analyst + Research Director

Published: April 2022

IN PARTNERSHIP WITH

VAST

IS THIS THE END OF TIERED STORAGE?

The requirements for data storage, processing, and compute have been fundamentally reshaped over the past two years.

Consider that the shift to digital has been swift, irreversible, and more extensive than most organizations ever anticipated. A majority of organizations today are still scrambling to adapt to new employee, operations, and customer engagement models.

As business models have been disrupted, new technologies, devices, and applications deployed both with and without the consent of IT organizations have shredded legacy tiered data storage and compute models that were driven by long-term IT plans, and not unanticipated, dynamic user demands.

Application modernization and cloud migration initiatives have also been accelerated, as have edge computing and IoT implementations. And let's not forget the ongoing development and adoption of digital technology that has been transforming businesses, brands, and organizations over the prior decades — technologies that have come to dominate previously non-digital industries such as the automotive and transportation sectors, manufacturing and operational systems, medical devices and research, and public utilities and services.

More digital means more data — a lot more data.

All of these changes have resulted in more, variable types of data being created — data that can be gathered, analyzed, and used to provide business insights and business value for those organizations that are agile and able to adapt. Unfortunately, many organizations are burdened with legacy tiered data storage models that just don't stack up with today's — or tomorrow's — demands.

For organizations looking to improve their data value, cost control is paramount as investments made today are often difficult to undo or replicate in the near future. The ability to scale is critical as the types, sources, and volume of data are changing rapidly due to new devices, new engagements, and new business models. Moreover, as use cases trend toward the real-time value of data, a new level of performance is required to generate the insights required to keep pace with customer and competitive demands.

The questions enterprises face today are important ones:

Is the traditional model of tiered data storage the right approach for today? Or is there a better, more effective, and more flexible model that offers greater opportunity at a lower cost?

THE FUTURUM PERSPECTIVE AT-A-GLANCE

Here are the key points and findings of this paper and our research into VAST Data's offerings:

1. Global disruption has become data storage disruption: New technologies, devices and applications deployed both with and without the consent of IT organizations, have shredded legacy tiered data storage and compute models driven by long-term IT plans, and not unanticipated, dynamic user demands.
2. The way we capture and store data has changed: Today's data is created, captured, and stored from digital systems distributed across the entire ecosystem, from the first supplier to the last customer, and from the core to the edge. It represents every stage of the design, production, and end-user experience and requires a highly scalable storage architecture.
3. The way we use data has changed: Traditionally, old data became historical data of limited value and was well-suited to tiered storage archives. In the data economy, however, old data ages much better and is accessed regularly to provide context and insights alongside real-time data. Data is the fuel of artificial intelligence, critical to the ongoing training of machine and deep-learning systems and must be always accessible.
4. The value proposition of tiered storage has passed: Designed for an era of scarce and high-priced technology with limited need for real-time, high-performance access to historical data, tiered storage solutions made economic sense for many organizations. But the same features that added value to tiered storage are new barriers to cost-effective scalability and limit the resilience and performance required by enterprises in today's digital economy.
5. The era of universal data storage is here: New technologies for the design and manufacturing of storage class memory, the development of disaggregated storage networking and fabrics, and the resulting ability to cost-effectively deploy and scale distributed storage solutions with the ease and performance of direct-attached storage offer the opportunity for organizations to untap the value of all data regardless of type, location, age, or application.

DATA INFRASTRUCTURE AND STORAGE CHALLENGES

Most data storage infrastructure installed over the past decade has used a tiered approach to matching data to storage technology based on several criteria, including how often the data is accessed, the level of speed or performance of access required, and currency (or age) of the data. This was in part an effort to control costs by only using newer, more expensive technologies for data that had that direct need while leveraging older, legacy (and less expensive) technologies for data that would be considered less hot and more cold (less frequently used).

New technologies, devices and applications deployed both with and without the consent of IT organizations, have shredded legacy tiered data storage and compute models driven by long-term IT plans, and not unanticipated, dynamic user demands.

For example, solid state drives (SSDs) were fairly expensive when they arrived on the market, but they did offer a vast improvement over older disk and tape technologies. Cost prohibitive to install and use as the only storage mechanism, SSDs were reserved for only that data that required SSD levels of performance. Other data was partitioned off into sub-layers of Network Attached Storage (NAS), data servers, and various backup or low-tech archival systems. All of this assumes that hot data, with premium tier requirements, will ultimately age and be migrated down the stack from high performance tiers to high capacity tiers.

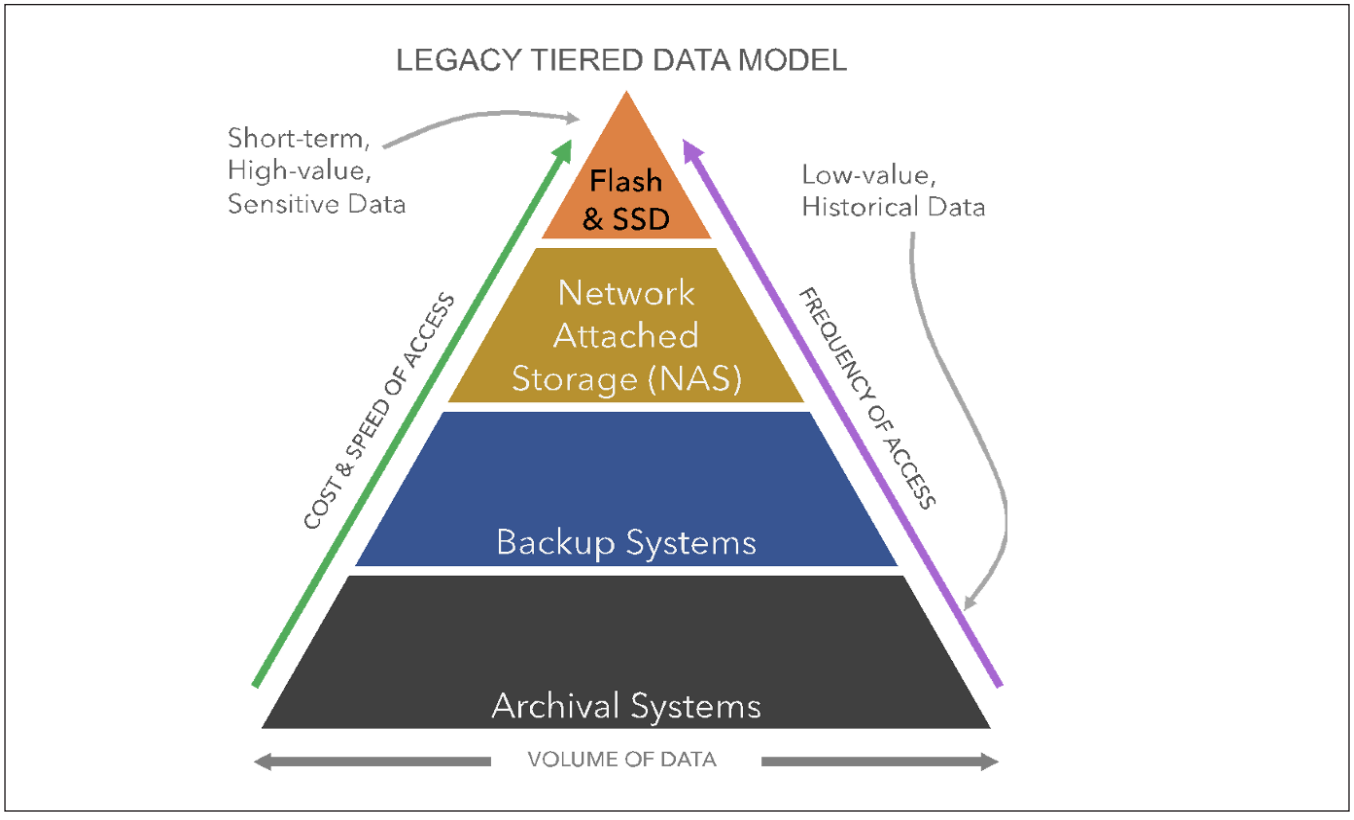

This model made sense in the days when storage technology was still focused on spinning media disk drives where incremental improvement often came with a steep price increase and a forklift implementation.

Original Tiered Storage Benefits

- Cost averaging, allowing enterprises to reduce the average cost/TB of their overall data infrastructure.
- Operational efficiency, as different types of storage resources are used only by the data that requires that level of performance, accessibility, or capacity.
- Legacy asset utilization, as investments in older generation technologies can continue to be used where appropriate to maximize prior investment value and limit new investment expenditures.

In today's environment, however, tiered storage imposes significant limitations on enterprise organizations and can, in fact, create negative value.

How we use data has fundamentally changed

In the past, data was accumulated to create an historical record or for compliance requirements — it was rarely accessed once archived and even then, from a select few.

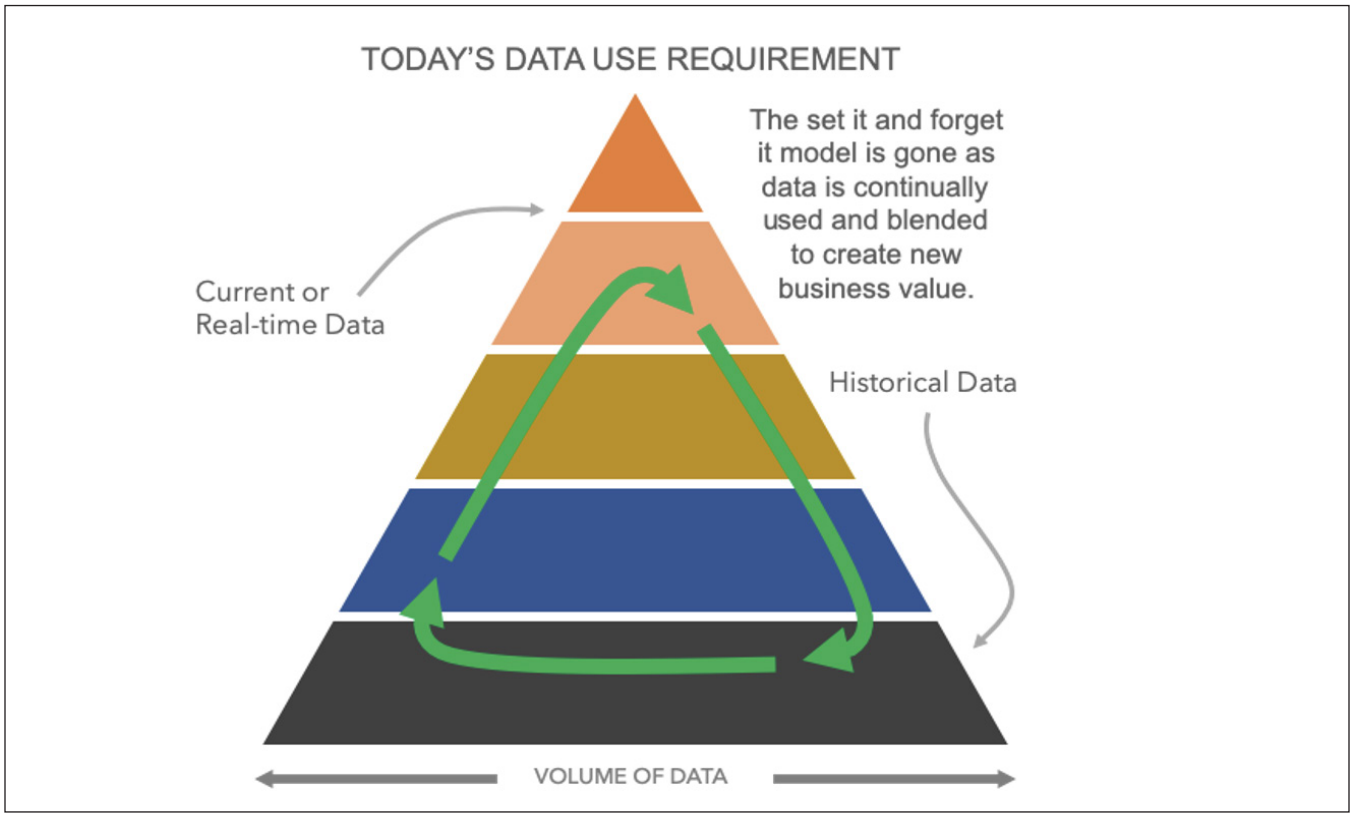

Today, however, historical data is regularly accessed by many to provide richness and context to current or real-time data — it's essential for today's AI applications, including the ongoing training of machine learning systems and the implementation of predictive analytics systems.

The tiered storage model no longer applies.

There are several key implications associated with the continued use of tiered storage models:

Management Complexity. Managing data in tiered storage is incredibly complex and many organizations don't have the support or capacity needed for this task. This is a significant issue for organizations that must deal with different technologies, different providers, and different operational processes associated with both different data storage tiers and the accumulation of different technology stacks due to acquisitions and prior independent implementations across different business units.

Siloed Data. Just like an object at rest tends to stay at rest, data within individual tiers tends to stay within those technology-limited boundaries. As a result, complex data movement, a result of data requirements that span multiple tiers, can be slow and result in significant application delays and prevent higher-performance applications from accessing data directly on lower tiers. And by limiting the data available to higher-level applications, the ability to make insightful decisions based on all the data can result in less than optimal decision-making.

Limited Scalability. Tiered data systems lack the ability to flex as a single system, instead requiring the highly coordinated provisioning — up or down — of individual tiers. This can also result in significant costs for technology, talent, and ongoing operational support.



ADDRESSING THE CLIENT NEED

As business needs change and evolve, so too do the requirements for how data is acquired, stored, and accessed.

The gap between current and aged data continues to shrink as organizations increasingly rely on historical data to guide real-time decisions for immersive experiences, customer engagements, operational efficiencies, and training AI and machine learning systems.

Organizations must find a way to implement the right storage solutions that allow their data to be utilized to the fullest, particularly as the volume and velocity of data continues to increase all while the value of aged data increases as well.

What is needed to meet the data storage demands of today and tomorrow?

Performance. Real-time data-driven insights are playing an increasingly vital role in businesses across all sectors. The speed of data access, including the ability to ingest and process vast amounts of data as it is created, is key to the ability to make decisions in the moment, where they matter the most. With the right level of performance that is not constrained by legacy storage technologies, businesses can step up to meet the real-time demands of providing efficient and meaningful customer engagements, supply chain and inventory management, improved asset utilization, predictive maintenance, and intelligent automation.

Scalability. Organizations are increasingly adopting technologies and methodologies that enable business agility. This requires the ability to cost effectively and rapidly scale up or down to meet changing business or customer requirements.

Unfettered Data Access. More data is better data when it comes to predictive analytics, machine learning, and deep learning. Unfettered, or universal, data access where all data is accessible all the time can provide significant business value, but that requires a much more consistent application of storage technology that treats all data as similar as possible.

SPOTLIGHT: VAST DATA's UNIVERSAL STORAGE

As discussed above, the tiered storage model originated in a time when it simply wasn't technically (or financially) feasible to implement a universally consistent data storage model. Today, however, the technology does exist and is at the core of the value proposition offered by VAST Data and its Universal Storage model.

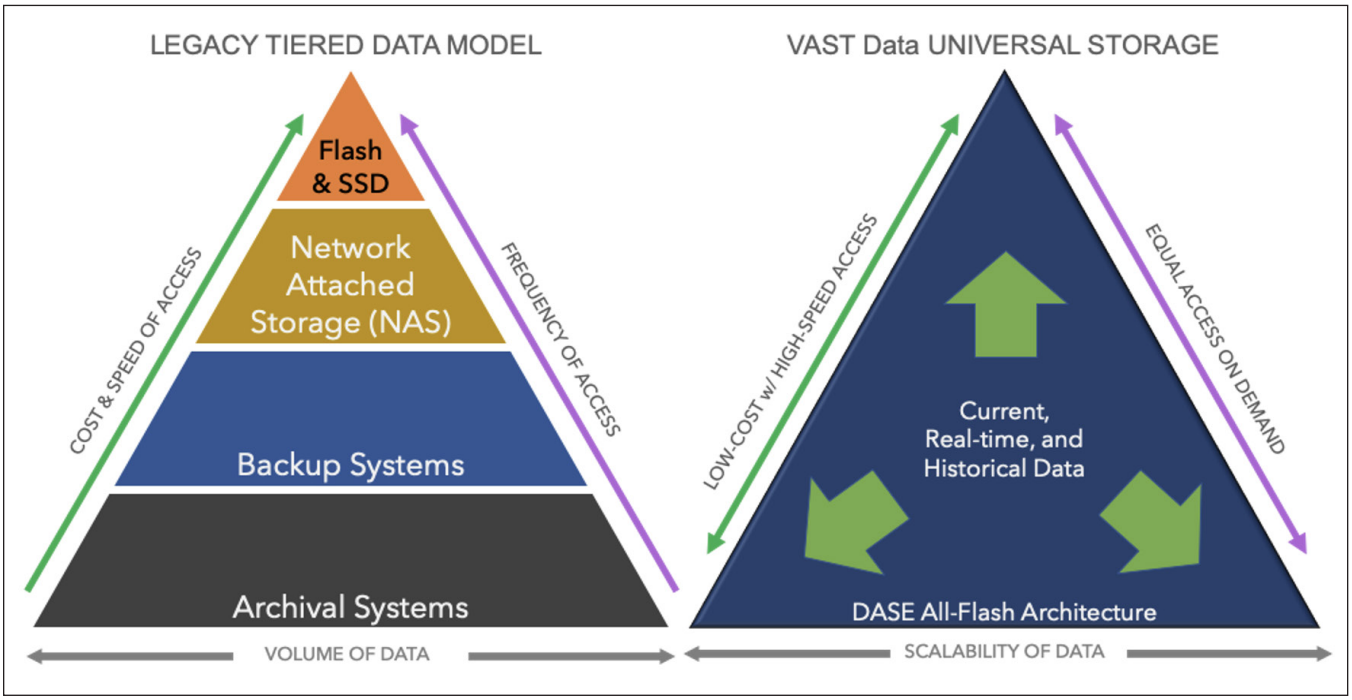

In observing VAST Data, it's clear they've embraced their mission to help organizations overcome decades of complexity and application bottlenecks by providing a simplified and universal storage experience. In contrast to a tiered model, VAST Data has developed its offering using a Disaggregated and Shared Everything (DASE) all-flash approach made possible by recent developments in Non-Volatile Memory express (NVMe) over Fabrics (NVMe-oF) and Storage Class Memory to take advantage of lower costs and increased resiliency of flash memory.

Key to VAST's value proposition are the algorithms and write-shaping software it has developed to minimize the wear on QLC and cost-optimize SCM. This IP is also a strong competitive moat against competition moving forward, an advantage we anticipate VAST will continue to develop as it expands its platform.

The essence of VAST's Universal Storage is one of simplicity — of technology, of design, and implementation. All data storage — regardless of where it may have resided in a legacy tiered model — leverage the same SSD solution, eliminating the complexities and limitations of tiered solutions. And while its storage systems are designed for exabyte-class hyperscale environments and all SSDs appear as local to protocol servers, they've priced the solution in line with a TCO below most tiered storage systems and well below what you'd expect for an all-SSD solution.

Flash Performance. Using both high-density QLC flash SSDs and non-volatile Storage Class Memory (SCM) SSDs and focused on file storage, not block storage, VAST has optimized its design to minimize read/write wear — a key component in VAST's 10-year support option.

Scalability. VAST's use of NVMe-oF enables the logical connection of remote NVMe devices into a single storage fabric that can be scaled as a single system without the need to worry about scaling individual tiered systems or rearchitecting server interconnects. Resources scaled up can be used by all data and applications, resulting in a much simpler approach to scalability.

Cost Efficiency. With a single-cost approach to memory that offers the performance of higher-priced top-tier storage components (yet is priced at mass market levels), VAST has developed a very cost-effective model that simplifies procurement. This includes its Gemini offering that allows customers to procure storage hardware at cost while adopting a flexible use-based managed software consumption model.

A key benefit of this type of simplified approach is in the forward-looking flexibility it provides. Complex, tiered systems are inherently difficult to both repurpose and build upon due to the specialized nature of their deployment. We believe VAST's customers can get the hardware cost efficiencies of a hyperscaler while preserving the simplicity of a software-defined business model.

VAST's Universal Storage improves the cost/performance of traditional data storage applications and is essential to the cost effectiveness and performance of emerging applications involving artificial intelligence (AI), the training of machine learning (ML) systems, and predictive analytics & maintenance solutions.

The elimination of complexity, the creation of a flexible use architecture, and support for multi-generational, asymmetric clusters all play well into what we expect will be a long-term data center infrastructure play as VAST builds on the competitive strength of its current offerings to reframe the traditional data center infrastructure model for the next generation of storage and compute requirements.

VAST DATA IN ACTION: Case Studies

The simplicity of VAST's Universal Storage solution allows it to drive operational efficiency and business value across a wide range of industries and use cases. VAST Data's customers span myriad industries across financial services, the public sector, healthcare and life sciences, media and entertainment, and manufacturing to name a few. Key logos the company works with include Jump Trading, Agoda, DUG, Squarepoint, Boston Children's Hospital, US Department of Defense, National Cancer Institute, Lawrence Livermore National Labs, Nanox.AI, Invitae, Translational Genomics, Verizon, Ingenuity Studios and many others. Here's a quick summary of actual case studies that demonstrate the breadth of VAST's value proposition.

CASE STUDY: Big Data and Analytics

Client: A provider of a travel sales internet platform.

Challenge: Improve scalability and move away proprietary DAS-based Cloudera HDFS and proprietary toolsets.

Solution: The organization adopted multi-protocol support for their diverse data ecosystem. By using NFS with RDMA and GPU-Direct Storage, they are able to power thair Apache tools with S3-compatible, fast object as well as support their Python and GPU-based machine learning code. All from a single unified space.

Value: The shift to a multi-protocol support environment improved scalability, simplicity, and flexibility in a complex data ecosystem.

CASE STUDY: AI-based Data Analytics in Medical Imaging

Client: A provider of medical imaging services using MRI and PET scan data for medical research.

Challenge: Improve the performance and scalability of its storage systems to allow AI deep-learning tools to leverage petabytes of image data to accelerate medical research.

Solution: A shift from legacy HDD storage to VAST's SSD-based Universal Storage system and disaggregated shared everything (DASE) architecture supporting NFS over RDMA (NFSoRDMA) fabric connecting stored flash data to NVIDIA DGX-A100 servers for training deep-learning applications.

Value: The shift to an all-flash storage model has allowed the firm to eliminate the complexity and management overhead of its previous HDD-based storage solution and scale storage and performance to the multiple-petabyte level required to train its deep-learning systems for image reconstruction, analysis, and parcellation. Operational cost savings have allowed it to redirect funds to additional GPUs to expand and accelerate research initiatives.

CASE STUDY: AI Analysis in Medical Imaging for Patient Diagnostics

Client: An AI analytics platform provider that allows healthcare institutions to better identify and diagnose disease for improved patient care.

Challenge: Needed low-cost object storage to build their deep learning environment.

Solution: A pre-GA version of universal storage was installed in the customer environment for deep learning workloads. The solution improved performance, resilience, and reliability.

Value: The shift to a universal storage solution has allowed the organization to stop fighting the storage economics, tiering, and complex integration systems and focus solely on training their AI models.

CASE STUDY: Petabyte-scale Backup and Recovery

Client: A global provider of data management platforms and data storage, protection, and recovery solutions for on-premises, hybrid, and multi-cloud environments.

Challenge: Improve the cost, performance, and storage requirements of its backup, restore, and recovery solutions.

Solution: Implement VAST's Similarity-based Data Reduction tools on top of VAST's Universal Storage platform increase compression, lower costs, and improve the speed of both backup and recovery. VAST Data currently partners with data protection companies including Commvault, Cohesity, Veeam and Rubrik to enable rapid backups and restores for instant recovery at a cost comparable to hard drive economics.

Value: A significant improvement in data reduction (compression) and the speed of backup and recovery operations.

- Compression improvements from 8:1 to 13:1 for VMware SQL Server Backups; 4:1 to 6:1 for Unstructured Data Backup; and 3:1 to 4.8:1 for VMware Backups
- Performance and scalability improvements for both Backup (from 18 TB/Hr to over 288 TB/Hr) and Restoration (from 144 TB/Hr to over 2,304 TB/Hr)

CASE STUDY: Using AI to Crack the Code on Cyberattacks

Client: A high-performance computer center.

Challenge: Cyberattacks continue to evolve and become more sophisticated. Protecting against the unknown can be challenging and cost prohibitive for organizations.

Solution: VAST Data was able to deliver affordable flash at petabyte and exabyte scale.

Value: Organizations are able to leverage the new hyperscale flash, storage class memory, and NIVIDA NVMe-over-Fabrics clusters to implement new algorithms for security.

CASE STUDY: Improved Performance in Genetics Analysis

Client: Genome analytics company providing cancer, IVF, and COVID mutation analysis.

Challenge: Scale data access and compute performance to improve full-genome mapping of 100-250GB per patient for the development of diagnosis and treatment plans.

Solution: Migrate storage and compute from a legacy system to a new architecture utilizing GPUs over VAST Data flash storage solutions.

Value: A 6x improvement in speed of data access and an even more impressive 30x performance in I/O operations/second (IOPS).

CASE STUDY: Real-time Analytics and Insights in Financial Markets

Client: Research firm providing data-driven analysis of global financial markets throughout North America, Europe, and Asia.

Challenge: Overcome traditional data storage and access performance limitations that restricted its ability to use emerging artificial intelligence (AI) and machine learning (ML) systems in the real-time analysis of global trading and market data.

Solution: Implement VAST's Universal Storage platform as the foundation to its on-premises High-Performance Computing (HPC) infrastructure.

Value: A 6x improvement in speed of data access and an even more impressive 30x performance in I/O operations/second (IOPS).

CASE STUDY: Scaling High Performance Computing as-a-Service

Client: A provider of HPCaaS (High Performance Computing as-a-Service) data analysis and modeling for applications requiring exascale computing.

Challenge: Improve the reliability, scalability, and cost-effectiveness of its HPCaaS offering by eliminating ongoing HDD failures, removing storage specific I/O application requirements, and extending performance to allow for true multi-tenancy.

Solution: Migrating to VAST's Universal Storage solution featuring Intel Optane SSDs, a disaggregated shared everything (DASE) architecture to isolate non-optimized I/O and limit multi-tenant performance impacts, and a scalable interconnect model using the NVMe over fabrics (NVMe-oF) protocol to provide direct-attached storage (DAS) performance and scalability in a distributed model.

Value: Leveraging VAST's data reduction and improved reliability, the company was able to significantly eliminate administration efforts and reduce costs by up to 85 percent compared to its HDD model – all while being able to offer exascale scalability to grow petabytes of storage as needed without any single point of failure. The firm has doubled its compute capabilities while leveraging VAST's remote appliance management services to reduce internal support costs.

CASE STUDY: Conquering Computer Vision in the Automotive Industry

Client: A global automotive manufacturer in the midst of shifting into the data economy and transforming both its digital infrastructure and its vehicles in support of always-connected and autonomous systems.

Challenge: Implementing a single data storage architecture and solution capable of spanning geographically distributed systems with a wide range of design, safety, operational, and anonymized in-vehicle data and computing requirements spanning backup and recovery to containerized big data applications.

Solution: The firm is implementing VAST's Universal Storage platform throughout its operations, including all new core and edge data centers.

Value: Beyond achieving the scalability and performance required to connect its distributed edge systems to each other and its core system, the firm has been able to substantial cost savings while leveraging VAST's approach to improve the capacity and availability of data from all sources for high-performance applications and improved automotive design processes.

THE NEXT STEP

The enterprise of today has more data than ever before. That will be true tomorrow, next week, and next year, as the flow of data from sensors, transactions, and devices is only going to increase with time. But so too will the potential value of that data, provided an organization has the right technology, solutions, and processes in place to capture, analyze and gain insights from all the data regardless of data type or age.

Finding the value in data, whether it's being used in the training of a machine learning solutions for predictive maintenance or to drive real-time decisions in an automobile or on the trading floor, requires an approach that makes no assumptions on the value of any particular data set and ensures that all data is treated equally in terms of accessibility, reliability, and scalability.

While tiered storage systems, designed for last-decades' data use models, may have some value they also inhibit the creation of new value and cannot offer the cost-effective scalability, resilience, and performance of today's storage class flash memory and disaggregated storage networking architectures.

Enterprises have a unique opportunity today to address the rapidly changing digital economy head on with solutions that offer:

- Flash Performance through SSDs featuring high-density QLC flash supported by non-volatile Storage Class Memory (SCM) for all data, not merely selected data;
- Scalability through the use of disaggregated architectures using NVMe-oF to logically connect remote NVMe devices into a single storage fabric that can be scaled as a single system without the need to worry about scaling individual "tiered" systems or rearchitecting server interconnects; and
- Cost Efficiency through the simplified procurement and management of flash memory that offers operational cost savings over HDDs and eliminates the need for higher-priced top-tier storage components.



Key Recommendations

We offer the following recommendations for organizations considering simplifying their storage investments:

1. Separate cost savings from operational gains. New storage class flash memory solutions can offer significant lifetime cost savings over traditional HDD systems (and even tiered SSD solutions) — that's a given. Instead, focus on operational performance, scalability, and the reduction of management overhead (talent/time) as key decision factors.

2. Match your current and emerging data needs and use models to the appropriate storage model. While new approaches to data storage can bring value to most any organization, it's important to preserve and leverage investments in legacy technologies if they provide value. Focus on those needs that can provide the most improvement in value first and migrate legacy data and systems only when appropriate.

3. Match your current and emerging data needs and use models to the appropriate storage model. While new approaches to data storage can bring value to most any organization, it's important to preserve and leverage investments in legacy technologies if they provide value. Focus on those needs that can provide the most improvement in value first and migrate legacy data and systems only when appropriate.



IMPORTANT INFORMATION ABOUT THIS REPORT

CONTRIBUTORS
Daniel Newman
Founding Partner + Principal Analyst
Futurum Research

Shelly Kramer
Founding Partner + Principal Analyst
Futurum Research

Fred McClimans
Senior Analyst + Research Director
Futurum Research

PUBLISHERS
Daniel Newman
Founding Partner + Principal Analyst
Futurum Research

Shelly Kramer
Founding Partner + Principal Analyst
Futurum Research

INQUIRIES
Contact us if you would like to discuss this report and Futurum Research will respond promptly.

CITATIONS
This paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title, and "Futurum Research." Non-press and non-analysts must receive prior written permission by Futurum Research for any citations.

LICENSING
This document, including any supporting materials, is owned by Futurum Research. This publication may not be reproduced, distributed, or shared in any form without the prior written permission of Futurum Research.

DISCLOSURES
Futurum Research provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

ABOUT VAST DATA
Our mission is to bring an end to decades of complexity and application bottlenecks. VAST combines a series of innovations to radically change the flash cost vs. capacity equation, democratizing the utility of flash for all data and all applications.

ABOUT FUTURUM RESEARCH
Futurum is an independent research, analysis, andadvisory firm, focused on digital innovation and market-disrupting technologies and trends. Every day our analysts, researchers, and advisors help business leaders from around the world anticipate tectonic shifts in their industries and leverage disruptive innovation to either gain or maintain a competitive advantage in their markets. Read our disclaimer statement here.

CONTACT INFORMATION

Futurum Research, LLC | futurumresearch.com | 817-480-3038 | info@futurumresearch.com
Twitter: @FuturumResearch

©2022 Futurum Research. Company and product names are used for informational purposes only and may be trademarks of their respective owners.