White Paper

High-Performance NFS Storage for HPC-AI

Sponsored by: VAST Data

Josephine Palencia Ashish Nadkarni
January 2023

IDC OPINION

Digital transformation is leading to rapid proliferation of use cases that rely on performance-intensive workloads. These workloads — which were once exclusive to academia and government institutions — are now going mainstream, thanks to an urgent need for gaining insights from accelerated analysis of diverse data sets.

A recent IDC survey found that more than half of organizations indicate data growth rates from 20% to 50% in 2022, and a third expect them to go as high as 99%. 80% of this data will be unstructured in nature with a heavy emphasis on iterative, time-to-value insights. Workloads that deliver such insights make use of algorithmic approaches and workflows such as modeling and simulation (also known as high-performance computing [HPC]), training and inferencing (used in artificial intelligence [AI] and machine learning [ML]), and (high-performance) data analytics. Traditional approaches to compute and storage — which are rooted in a general-purpose infrastructure architecture — fall short of the performance and scalability demands of such approaches.

Investments in such workloads require accompanying expenditure on purpose-built compute, storage, and networking infrastructure that can host one or more of these performance-oriented environments. This infrastructure stack is inherently complex and highly involved with many closely interconnected parts that cannot easily be disaggregated from each other. With the convergence of these workflows and approaches, there is a growing requirement that the storage infrastructure support the divergent I/O patterns of traditional HPC modeling and simulation (small data in/large data out) and emerging AI/ML training and inference patterns (large data in/small data out) simultaneously. Historically, these have differed leading to a bifurcation in the storage infrastructure. This has led to higher costs and runs counter to the economics of digital transformation.

Emerging vendors such as VAST Data — an enterprise provider of scale-out unstructured data storage solutions — are seeking to change the situation. VAST Data, for example, offers a unique architecture to meet these requirements for both enterprises and institutional customers seeking to deploy a converged environment for traditional HPC and emerging AI/ML workloads. Adoption of this solution — which uses high-performance Network File System (NFS) for data access — leads to better agility, performance, and scaling of the converged HPC-AI infrastructure environment. It also makes the environment much simpler to manage, with enterprise-grade data services that are also designed to fit the needs of the institutional users. IDC found that VAST Data's solutions are being adopted in financial services, government, healthcare, life sciences, media and entertainment, manufacturing, telecommunications, and education markets, among many others.

SITUATION OVERVIEW

The Journey So Far — Why the HPC Community Moved Away from NFS

In the early days of modeling and simulation deployments (what we refer to as traditional HPC environments), the HPC community wholeheartedly endorsed the Network File System as an access protocol for data sharing across the HPC clusters. NFS had gained popularity for its wide adaptability, versatility, and ease of use. As a tried and tested IP-based protocol, NFS provided a cost-effective solution for network file sharing across multiple heterogeneous compute nodes. NFS further simplified file management, eliminated storage silos, reduced file transfers between nodes, and provided transparent file tiering capabilities.

Despite these advantages, NFS presented inherent scaling challenges in parallel computing environments where performance trumps everything else. These were addressed via high-bandwidth, low-latency networks and updated versions of the protocol, but the basic protocol-level scaling challenges remained. For example, if multiple clients write to a cached copy of the same file on the network and issue a close command in proximity, the protocol cannot ensure which bits become the authoritative version of data, leading to data inconsistency. It is the traditional implementations of NFS-based systems that introduced the problems for scale out.

As traditional implementations of NFS-based systems were never designed for atomic parallel I/O, many practitioners in the HPC community shifted to other parallel (distributed) file systems such as Lustre that were specifically designed for parallel, clustered computing environments. The parallel file systems have their own driver that needs to be maintained in line with the software versions running on the server side used separately from the NFS driver. These environments make use of a file system-specific client module — installed on each of the compute nodes — that augments the stock NFS driver for shared data access.

This shift was profound. Even to this day, most HPC environments — and especially research institutions — rely on parallel file systems. IDC finds that for organizations that deploy these parallel file systems, a day or few days of cluster downtime is considered an acceptable risk and enterprise features are considered less relevant.

For example, at the highest tier of the world's fastest supercomputers, Lustre remains widely adopted with 40% share, followed by other commercial and open source file systems. Frontier, the world's first exaflop supercomputer deployed at the U.S. Department of Energy's Oak Ridge National Laboratory (ORNL), uses a three-tier Lustre file system with 679PB (disk based) at 5.5TBps read, peak 4.6TBps write, 2 million random read IOPS; flash-based NVMe devices for 11.4PB at peak 10TBps read/write, 2 million random read IOPS; and flash-based metadata capacity of 10PB.

Challenges with Parallel File Systems

Parallel file systems typically extend the POSIX caching semantics, which are defined only for a single host (where they are much easier to be implemented efficiently) to a cluster of hosts (where the corresponding synchronization effort typically slows things down, to a degree where things like shared file writes become unfeasible and additional intermediate layers must be added). Published studies in the public domain indicate that parallel file systems, such as those that support burst buffers, often perform poorly for read-heavy random access workloads.

Further, the shift to parallel file systems itself introduces higher operations challenges into these HPC environments:

- **Skilled staff.** With the shift to parallel file systems, it became mandatory to have highly skilled HPC systems administrators to manage and optimize the setup. Even today, these skills continue to be in dire shortage in the workplace, with large service providers realizing the talents of such storage administrators and hiring them away, causing a cyclic requirement to train and cycle through talent.
- **Version control.** Parallel file systems need to maintain a certain consistency with kernel versions and OS to have compatible versions on clients and servers. Because parallel file systems do not use a standard protocol, their proprietary version can change with every update.
- **Implementation.** Parallel file systems require meticulous planning during the implementation phase. Many of the tuning parameters cause tradespace between latency and throughput that cannot be addressed given the architecture of the systems (e.g., setting up default inode maximum limits, metadata management, tuning parameters for known I/O striping patterns for applications, enabling HA, and setting up failover parameters).
- **Change management.** Offline and online configuration changes can often upset the delicate balance for optimized performance. Any configuration modifications on servers, volumes, and targets or relabeling of device names during multipath configuration outside an integrated software manager heightens the risks of breaking HA, file system management, and monitoring capabilities. In addition, many of the updates to and expansion of the system require downtime as they cannot be performed while the environment is operational.
- **Networking complexity.** In a heterogeneous network environment using mixed generations of InfiniBand and Ethernet networks, gateways and routers need to be incorporated, further introducing configuration management and higher management overhead.

Operational overhead in turn translates to higher costs. When the total cost of the implementation is considered, this can make a significant difference.

Making the Choice for Storage Systems with High-Performance NFS Access

With HPC and AI use cases converging, the overall nature of workloads is changing. Fortunately, the technology landscape is evolving in ways that when designed with the right architecture, storage systems (with NFS-based access) could provide the scaling and agility required for HPC-AI environments.

Overcoming the Perception of NFS in HPC-AI Environments

The stigma and hesitation associated with using NFS as the primary access protocol can be reduced with the right scale-out architecture for the underlying file system and protocol. It requires addressing parallelism in the storage itself rather than relying on the protocol, just leading to significant improvements in performance and scaling characteristics.

The good news is that in the HPC community, NFS is still very much in use. In traditional HPC environments, NFS-based storage is often delegated as a protocol for storage access across a heterogenous and hybrid infrastructure. In fact, IDC's research finds that most HPC environments use as many as three storage tiers (primary, secondary, and tertiary) — with NFS often used as the secondary or tertiary tier. This tiering categorically introduces higher technical and management complexities, data silos, data fragmentation, and unnecessary data migrations.

More recently, IDC is also noticing a shift among more mature HPC research organizations. Many seek to benefit from the features and ease of use of enterprise storage systems. With pressure to improve service quality, they find that features like snapshots provide significant additional business value to easily recover from accidental deletion of important files.

However, they struggle to find solutions that are easy to use and full featured. Crucially, they require storage solutions that can provide the scale-out performance of a parallel file system at an affordable price point.

NFS Is Much More Mature Today Than a Few Years Ago

Fortunately, the lack of adoption of NFS as a primary protocol in HPC has not deterred vendors and developers from continuing to invest in NFS. The rapid growth of the open source ecosystem led by Linux has led to the expeditious development of the network-attached storage (NAS) industry; growth of high-bandwidth, low-latency TCP/IP networks; and applications that are certified to work with shared network-attached file systems. The main beneficiaries of the maturity of NFS have been enterprises. Growth of unstructured and semistructured data has led NFS to be widely adopted in organizations in most industries.

Growth of AI/ML workloads in the enterprise is providing further boost to NFS adoption. Several interviews with IT practitioners implementing AI/ML workloads in the enterprise over the last year reveal that most enterprises have some type of highly scalable file system with shared networked access, usually via NFS. For most workloads with modest I/O requirements, NFS works well, though once it starts to get into scale it can run into issues. Common requirements that surface these performance problems include a need for extremely high sequential read/write performance to a single large file and the ability to scale to the needed levels of concurrency as additional clients and applications accessing the data are added.

Converged Infrastructure for HPC-AI Reinforces the Need for NFS

Different approaches have ruled how enterprise IT and traditional HPC environments have been deployed for a few decades now; in the enterprise, it is the traditional three-tier architecture, while in the HPC community, it is a software-defined, massively scale-out, or parallel architecture. The convergence of HPC, AI/ML, and high-performance data analytics (HPDA) use cases means that most enterprises and institutions can pursue a common infrastructure architecture that makes use of the best of both worlds.

NFS-based approaches warrant a second look in the case of the storage layer for such an architecture. NAS systems are designed to scale with the use of the underlying file systems, a massive namespace, and scaling capabilities that rely on high-bandwidth low-latency networks. With the right parallel architecture, such systems can scale to hundreds of petabytes.

IDC believes that combining the benefits of NFS, flash media, and parallel scale-out architecture on the server side provides maximum benefits for HPC-AI environments. Such a system enables a common infrastructure layer for traditional HPC as well as converged HPC-AI workloads. It dramatically reduces deployment complexity and administrative burden.

Changing HPC-AI Spending and Use Cases

As HPC goes mainstream and converges with AI/ML, data analytics, and engineering workloads, the market is witnessing a shift in the demographic of the demand. HPC has expanded its footprint into all organizations undergoing digital transformation seeking to establish optimal performance and efficiency in their infrastructure on premises, in the cloud, and at the edge. As a validation, IDC's 2022 *Worldwide HPC Survey* received a mix of respondents from corporate businesses, for-profit institutions (e.g., business start-ups and partnerships), traditional clients, national and government research laboratories, and academic institutions (e.g., universities).

In terms of spend, IDC finds that in 2022, most organizations in North America (the United States), Asia (China and Japan), and Europe (France, Germany, and the United Kingdom) indicated that they will spend a third of their total HPC-AI budget on high-performance storage.

For storage vendors, this presents an excellent opportunity to highlight and differentiate in a small storage market with exceptional storage performance reinforced by a comprehensive portfolio that supports the entire data logistics and data management operations (e.g., policy engines, data protection, data classification, metadata management, replication services, disaster recovery, and redundant backup). Additional considerations include HA; a high degree of flexibility for integration in a heterogeneous HPC-AI compute, storage, and network environment; business agility across all platforms (on premises, cloud, and edge), ease of use; and, most importantly, the short- and long-term ROI/TCO analysis for cost-saving targets of CEOs.

VAST PROPOSITION FOR HPC-AI WORKLOADS

VAST Data is an enterprise provider of scale-out unstructured data storage solutions. The company offers a unique architecture to meet these requirements for both enterprises and institutional organizations looking to deploy storage for traditional high-performance computing environments, emerging AI/ML workloads, and even a common converged stack for HPC-AI environments. In doing so, VAST Data has quickly gained the stature of a trusted provider for bellwether names in the financial services, government, healthcare, life sciences, media and entertainment, manufacturing, telecommunications, and education markets, among many others. The vendor counts three of the top cloud service providers and eight of the Fortune 50 among its customers, a fact that speaks to the economics of its platform at scale.

VAST systems are now deployed in many top 500 datacenters. The VAST Disaggregated Shared Everything (DASE) architecture is designed to scale to meet the most demanding infrastructure requirements of the HPC-AI era (to over 10,000 storage servers and 1,000 NVMe enclosures totaling 1.5EB scaled to 100TBps with over 500M IOPS).

The VAST Architecture

To meet the requirements of digitally transforming organizations, VAST Data sells an all-flash scale-out storage system — dubbed Universal Storage — that is based on commodity x86 server-based storage hardware coupled with a software-defined storage controller platform.

The base hardware platform of the VAST Data Universal Storage consists of optimized NVMe enclosures that range from 338TB to over 1PB of capacity per enclosure, storage-class memory (SCM) of 18/6.4TB, performance of 40/64GBps, and network options of 4x/8x (100GBE/EDR IB). The 4x Stateless VAST servers/chassis has 80 x 2.4GHz cores per chassis, 1TB DRAM per chassis, and 8x bandwidth (50/100GBE and/or IB).

VAST's Disaggregated Shared Everything architecture uses a software stack that applies global algorithms for peak performance and scale. The state and storage capacity of the system are built from high-density NVMe-oF storage enclosures, and the logic of the system is implemented by stateless Docker containers that connect and manage all media. Since the compute elements are disaggregated from the media across a datacenter fabric, each can scale independently, decoupling capacity and performance and enabling them to be scaled independent of each other. By eliminating the need for intracluster coordination and by scaling across high-throughput commodity networks, the NFS file services can scale to exascale proportions across a distributed network of attached storage arrays. According to VAST, DASE is the first new data architecture since the introduction of the shared nothing concept in 2003.

The comprehensive VAST Cluster software stack supports NFS, NFSv4 with Kerberos authentication, NFS over RDMA, SMB, and Docker CSI and S3 access methods and leverages newer storage technologies like storage-class memory, NVMe over Fabrics, and high-density quad-level cell (QLC)-based solid state disks (SSDs). All VAST Data customers use it as a scale-out file system, but over 40% also use it for S3 access.

VAST Data supports a variety of enterprise-class storage management features that enable high availability, dense multitenant workload consolidation, and extremely efficient utilization of storage compute and capacity resources. The vendor delivers an enterprise-class, unified data platform that enables dense consolidation of a wide range of workloads without risking performance and/or availability service-level agreements (SLAs) that is far simpler to deploy, manage, and scale than more traditional scale-out storage designs. For data services, the system supports an exabyte-scale namespace, snapshot and object immutability, distributed file locking, LDAP and Active Directory integration, global data compression, multitenancy, multiprotocol access to the same data, access control lists, capacity and performance analytics, and zero-impact snapshots.

Finally, the advantage of VAST over parallel file systems is the former uses NFS as a standard protocol that does not change with upgrades. Hence when VAST servers get upgraded to new versions, it does not require corresponding changes on the clients. This eliminates the necessity for considerable planning and implementation during upgrades that can entail operational disruption and/or downtime.

The VAST architecture is illustrated in Figure 1.

FIGURE 1

VAST Architecture

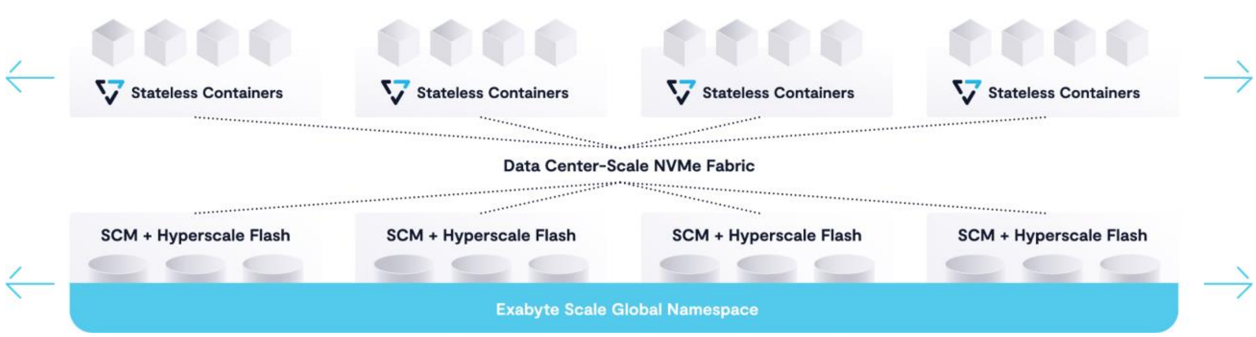

Source: VAST Data, 2023

Built-In RDMA Support

The Universal Storage platform can be connected over Ethernet and InfiniBand networks without the need for gateways or protocol routers, thus providing flexibility in multivariate networks. With the addition of RDMA, NFS performance can go beyond the limits of TCP/IP. The system can be accessed via either NFS or NFS over RDMA thus enabling a single NFS client to saturate a 100Gb connection (e.g., 11GBps over EDR). VAST servers can be pooled into resource groups that enable the cluster to be accessed via multiple network subnets without needing each subnet to have access to the entire server pool. This enables complete network disaggregation, permitting each server to have access to the global state of the cluster over NVMe-oF.

Global File Locking, Cache Coherence

The VAST Universal Storage platform supports Network Lock Manager (NLM) byte-range locking protocol to request and release locks on files and byte ranges. NLM allows shared and exclusive locks to applications and is designed for parallel applications where many byte ranges can be locked concurrently within a single file. VAST Data's approach to NLM locking is inherently scalable because the locking and lock management are fully distributed across the VAST Cluster. VAST Clusters leverage the DASE architecture to eliminate the need for centralized lock management. Instead, lock information is stored as extended file system metadata for each file in the VAST V-Tree, distributed globally across the system's SCM. Since all system metadata is available to all the storage nodes in the cluster, each storage node can create, release, or query the lock state of each file it is accessing without the central lock manager server that can become a bottleneck on systems. While optimized for random I/O where data is first written to the SCM, there is the option to disable client caching when not required by the application. When disabling client caching by using the sync or o_direct options, writes become atomic operations and the client consistency issue in legacy NFS systems is resolved. Applications that require caches turned on use global file locking.

Global Data Compression

VAST Data's innovations in data reduction (known as Similarity-based data reduction) combine compression and deduplication into a new capability that dramatically reduces file and object data footprint. The compressor scans the data for repeating patterns, replaces them with smaller symbols, and builds a dictionary so the decompressor can reverse the process. It works globally across all files (block-based deduplication) and the pattern matching is byte granular. Since the pattern matching works at 1/4,000th the granularity of traditional deduplication algorithms, it is less sensitive to data differences. This feature reduces data that has already been compressed by applications. In many large-scale HPC environments, 2.5:1 is a commonly observed data reduction ratio. In practice, data reduction has been achieved with pre-compressed data in genomics (2:1), seismic data (2.5:1), backup appliances (3.1, pre-duplicated), animation/VFX files (3:1), index and log files (4:1), quantitative trading market data (8:1), and backup data (20:1). These reduction ratios are for live data, and since VAST data reduction is always-on, they reduce all data at all times.

VAST Data's All-Flash Cost Remediation

The Universal Storage platform is designed as a single-tier, all flash-based storage system that provides the economics of HDD-based storage. The all-flash storage strategy is more important now than ever as the HPC (write-oriented) and AI (read-oriented) workloads converge onto a common performance-intensive computing infrastructure that consists of a mix of processors and coprocessors/accelerators (e.g., CPUs, GPUs, and FPGAs).

The read I/O is also characterized by random, lots of small files (LOSF) that cannot be deterministically prefetched into some buffer or cache layer in a storage. Cache misses that go down to HDD storage can cost organizations as much as a 98% loss in performance and impact AI computing efficiency. This is because HDDs cannot manage random reads well. To read a random block on a disk, the head needs to move to the right track along the radius of the platter (seek time up to ~15ms). Then the platter needs to rotate until the right sector or block is under the read/write head (rotational latency ~6ms). For this reason, the HDD at 4K random reads performs at a dismal 1–3MBps at 50–90 IOPS.

VAST Data's all-flash approach enables the company to service both HPC and AI clients with high-density throughput as well as limitless IOPS with the price point of a less performing HDD solution. As for cost remediation for using expensive all-flash components, VAST Data constantly and aggressively optimizes hardware and the software layers:

- **For hardware:** Efficiency and performance are gained by integrating NVMe-oF, QLC (4 bits/cell), next-generation PLC (5 bits/cell), and SCM.
- **For software:** Data compression and reduction techniques with similarity hashes and adaptive chunking with variable block length provide 2–3x benefits relative to similar products.

Overall, these effectively preserve and increase storage capacity and significantly reduce the costs for adopting a VAST all-flash system.

VAST Data has targeted three of the goals most enterprises and institutions are seeking to achieve when doing a technology refresh:

1. Moving to a more efficient, agile, and scalable infrastructure architecture
2. Consolidating performance-critical workloads like HPC and AI/ML
3. Reducing operating overhead and costs

While established enterprise storage providers have tried to evolve their existing architectures in a direction better suited to the requirements of this new era, a few start-ups have started with a blank sheet of paper just within the past several years to design a storage architecture from the ground up to meet evolving storage needs. In the area of unstructured data storage infrastructure, VAST Data has indisputably distinguished itself with a unique converged storage proposition.

IDC interviewed VAST Data customers from two industries (academic research institutions and financial services) to gain a perspective on the all-round adaptability of the Universal Storage platform.

Leading Research University

Use Cases: Genetics, Genomics, Modeling, and Simulations: Computational Fluid Dynamics (CFD), High Energy Physics

IDC interviewed the head of Research Computing at a private research university with oversight of two large HPC environments. The IT team consists of five staff members supporting 2,200 users.

The HPC infrastructure is made up of 2,000 servers with approximately 10,000 processor cores, 200 GPUs, and 2PB of VAST Data. The university has transitioned through several generations of GPUs and presently reports GPU hour usage in the billions (1/150 CPU). The HPC environment is largely based on an Ethernet fabric with some InfiniBand, but the university does eventually plan to transition to all Ethernet. The software stack is all open source based, which is typical of most universities.

The university installed VAST less than two years ago ahead of the enormous demand it has seen for its HPC environments. As a result, it plans to double its entire infrastructure in a few months. The decision to go with VAST Data — after evaluating other storage systems in the market — was made on the basis of management overhead and the ability to scale performance as the environments grew in size.

The university's experience with VAST has been very positive. The environment has had zero downtime thus far thereby improving service quality tremendously. The university's IT administrators — who are also researchers — are happy with the ease of management and maintenance (a few hours of maintenance every quarter) and fast turnaround for support (resolved in minutes). High integrability (NFS, SMB, and no storage tiering) has allowed the university to rapidly ramp up its storage infrastructure without dealing with the HPC storage complexity. It has enabled the IT staff to devote more time on their research.

As early VAST adopters, the IT team believes they have not yet reached the limits of their HPC environment in terms of performance, throughput, latency, and scale. Their experience with VAST has given the staff the confidence to transition to use the Universal Storage platform as the single and only HPC storage solution for all their applications within six months. "Every university I know already has VAST installed (referring to six other universities)," said the interviewee.

Financial Trading Firm

Use Case: AI/ML/Deep Learning Training and Inference

IDC interviewed the director for Financial Services at a leading financial services firm with oversight of four HPC environments supporting global trading operations. These environments are located in the New York, Singapore, London, and Norway. The firm's all-on-premises HPC infrastructure consists of clusters with 2,000 servers and 10PB of VAST Data across all locations. The firm employs three IT administrators to manage the storage infrastructure. Prior to choosing VAST Universal Storage, the firm investigated many on-premises and cloud-based storage solutions. Ultimately, VAST's solution won thanks to its ease of management and maintenance, enterprise-grade data services and, crucially, performance at scale.

Deploying VAST was part of the firm's larger IT digital transformation strategy to streamline infrastructure and simplify support while delivering performance and scalability for the HPC environments. The firm has seen tremendous business value in transitioning to VAST:

- The VAST environment has made IT operations more agile and responsive to business demands by being able to run more jobs without additional hardware. For example, the firm's research stack has 10,000 concurrent jobs that can be easily managed on the VAST platform. The time needed to deploy new storage resources has improved by 90% (to five hours from an average of two days).
- After the transition to VAST, the firm has not experienced any slowdowns or outages for research or production jobs. For example, the number of monthly unplanned downtime incidents per month is down to zero from an average of three, while the monthly average time for unplanned downtime is down to zero from an average of six hours. With no downtime, there is no loss of user productivity, no data loss, and no impact on the business.
- The company found the VAST Data Universal Storage platform to be comparable costwise to HDD-based enterprise storage solutions. The firm achieved additional operational cost savings with VAST through better support and zero downtime and by switching to NFS (thus forgoing purchase of proprietary storage licenses).

ESSENTIAL GUIDANCE FOR HPC-AI IT DECISION MAKERS

Any customer looking to deploy a new file- and/or object-based storage platform that will start near 1PB of effective capacity should evaluate VAST Data's Universal Storage. Given the growth of most data systems, the need to scale out storage is increasingly important. In addition, the confluence of AI/ML and HPC has led to a requirement for analytics across the entire corpus of data sets, which are limited by the scalability of most classical enterprise architectures. Workload mixes that are not expected to grow to 1PB and beyond and do not benefit from any of the features of solid state storage may not enjoy a good economic value proposition relative to HDD-based systems — the system's unique architecture and features deliver its best value proposition at scales of 1PB and beyond as systems require larger data sets for AI/ML workflows. While VAST Data does offer an entry capacity point of 100TB, VAST Data's revolutionary solid state economics do not really kick in at that level, and potential customers looking at this low-entry capacity may only want to consider it for a very rapidly growing data set that will get them to 1PB or beyond.

Given VAST Data's ability to deliver submillisecond latencies across mixed workloads at scale, the Universal Storage platform offers the opportunity to effectively run a wide range of transactional, analytics, and other workloads on top of unstructured data. The fact that the data is accessible through multiple access protocols can facilitate data sharing among workloads, a feature that can be of significant interest in big data and analytics workloads. Other features of the system support other characteristics that are of interest with transactional workloads as well (atomic writes, extremely rapid recovery in the event of failures, etc.). A unique characteristic of the Universal Storage platform is its ability to meet the requirements of the metadata-intensive mission-critical workloads at the same time as meeting the throughput, scalability, and cost requirements of more traditional scale-out workloads.

OPPORTUNITIES AND CHALLENGES FOR VAST DATA

VAST Data has seen unprecedented growth in adoption across industries showing versatility, wide adaptability, and scope. IDC recommends VAST to continue building momentum across the performance-intensive computing market showcasing the might of the Universal Storage platform across a variety of use cases and industries as well as provide performance benchmarks, which are standard in the HPC community. A few examples to showcase the opportunities ahead of VAST Data are discussed in the sections that follow.

Federal Government

VAST Data Federal, a subsidiary of VAST Data focused on federal agencies, partnered with the Department of Defense (DoD) on a $10 million contract (2021) to modernize the agency's data-centric operations with an all-flash system to respond to national security threats in real time. This includes the Synthetic Agency Radar (SAR) project that creates 3D models and imaging of landscapes and buildings that channel massive amounts of data through VAST. The Lawrence Livermore National Laboratory (LLNL) adopted VAST for therapeutics research with a billion compounds to improve drug efficacy. Through this method, it was able to fuse five LLNL supercomputers to use VAST.

Telecommunications

Verizon and NVIDIA linked with VAST to manage volumetric video by leveraging NVIDIA's Magnum IO for parallel intelligent datacenter to support Verizon's 5G. Recording live action volumetric video requires 50–100 Ultra HD cameras to capture the action from every angle and encoded to generate immersive displays. This requires large GPU compute and generates terabytes of data; the GPUDirect RDMA can capture 250TB of volumetric data every half hour to make this feat possible.

Financial Services

Many financial firms in the highly competitive world of algorithmic trading adopted VAST for their storage. For example, Jump Trading, a global trading firm with 20 offices in the United States, Europe, and Asia, relies on fast access to massive amounts of data. Other adopters include hedge funds involved in quant trading.

Healthcare

The Broad Institute (MIT/Harvard) allied with VAST to expedite its frontier research on genomic medicine with AI-driven gene sequencing. The Boston Children's Hospital, ranked number 1 in the nation with worldwide collaborations focused on childhood disease cohorts and genetics/genomics research, streamlined processing of cryo-EM data collected within two hours using VAST ~400TB storage for acquisition, harmonization, integration, fast access, rapid analysis, and sharing.

For VAST Data, it finds itself at a critical point in the technology industry. High-performance storage has been one of the remaining last vestiges of an unsolved dilemma in the HPC community. The next big challenge is to pave the way for exabyte, zettabyte, and yottabyte storage to become commonplace in HPC-AI circles. Controlling storage costs while scaling up to exascale performance levels and beyond will continue to be challenging and will require rigorous testing of exabyte storage. This is a welcome challenge for VAST Data as its charts outs its future with a unique offering developed through decades of experience. VAST stands out and is highly differentiated in the small but highly selective HPC community. Displacing decades-old incumbents may still take time, and VAST Data's unique NFS storage solution will continue to require continuous and ongoing customer education to reshape prior assumptions.

CONCLUSION

The holy grail for HPC storage is a single solution that manages the entire spectrum of applications including classical HPC modeling and simulations, engineering applications, AI/ML/deep learning training and inference, and forthcoming quantum computing applications. VAST Data Universal Storage has all the necessary building blocks and approach to scale to exascale performance while providing the much-needed simplification to a complex, heterogeneous, hybrid infrastructure on premises, on the cloud, and at the edge. VAST can also continue to reinforce and expand its offering with a comprehensive, sophisticated portfolio detailing support for the entire data logistics and data management operations (e.g., policy engines, data protection, data classification, metadata management, replication services, disaster recovery, and redundant backup). As the HPC and AI community continues to survey available high-performance storage solutions that are cost effective, easy to implement, and easy to integrate and improve the organization's overall business agility in the fast-pace and accelerating technology race, VAST's NFS-based approach to HPC-AI presents a strong and most compelling case for consideration.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

140 Kendrick Street
Building B
Needham, MA 02494
USA
508.872.8200
Twitter: @IDC
blogs.idc.com
www.idc.com