

OVERCOMING
THE REAL
ROADBLOCK
TO THE
ALL-FLASH
DATA CENTER



All-Flash storage is taking over many, but not all, production storage responsibilities but because of cost, most organizations can't justify placing all production workloads on all-flash storage. These organizations then certainly can't justify placing secondary storage and data protection workloads on all-flash. These systems, in most cases, have the performance capabilities to support all of these various workloads but the high cost of flash relative to high capacity hard disk drives makes the goal of an all-flash data center impractical.

If the cost variable of all-flash systems can be resolved without sacrificing performance then most data centers would happily convert to all-flash. The system would also need to scale from a capacity standpoint to meet the demands of storing all of the organization's data. Since most scale-out systems, scale compute and capacity in lock-step, it is likely that most organizations would have more available storage performance than they would need.

Most importantly, a single all-flash system that stores all data simplifies storage management significantly. Today, IT administrators often have five to six primary storage systems, two or three secondary storage systems and two or three backup storage systems. They not only have to manually decide what data should be on which storage system but they also have to learn all the various management interfaces for each storage system.

The elephant in the room of course is how can the cost of all-flash storage be driven down to the point that it is competitive with hard disk drives and maybe eventually tape technology? The answer is multi-faceted including rethinking media protection (erasure coding) and data reduction algorithms. It also means effectively using Quad-level Cell (QLC) flash which has a price profile that compares favorably to high capacity hard drives. The problem is that QLC has a relatively low endurance threshold and most data center class storage system vendors have deemed it a consumer only media.

VAST Data, promises to solve the QLC problem through a combination of technologies to make the concept of a single all-flash buffer a reality for data centers.



VAST's single storage system with the performance of flash, convenience of NAS and the scale of an object store simplifies data management while improving organizations ability to extract value from their data.



INTRODUCING **VAST DATA**

VAST Data launched in 2019 with \$80 million of funding.

They began selling their solution in late 2018 and, at launch, had several customer purchases of over one million dollars. The goal of their storage solution is to **deliver to customers a system that can store all of their data, regardless of classification, on flash storage.**

VAST has built a universal storage system that addresses every aspect of the typical data center by delivering all-flash performance and Buffer-5 cost efficiencies. The product is designed to deliver exabyte scale and uses NFS and S3 protocols as well as extensions for RDMA.

The VAST solution counts on the reality that most customers don't need the extreme performance that higher end systems claim and that most customers are already using a file-system for most of their data and anything that doesn't, can be converted to run on a file system.

The **VAST Data solution** counts on three new technologies:

1

The first is NVMe over Fabrics, which VAST uses for compute to storage-shelf interconnect, resulting in each node having very low latency access to storage media. The high speed network enables VAST to leverage loosely coupled, stateless compute nodes which connect directly to the storage enclosures. The cache-less, stateless design eliminates the node inter-communication that becomes a scaling problem for typical scale-out storage systems. Each enclosure can provide 18TBs of storage class memory and 675TB (raw) of QLC capacity. Customers can scale performance and capacity independently.

2

The second technology is SCM/Optane/3D XPoint as a primary metadata store and buffer. The storage class memory buffer enables VAST to deliver incredibly high performance to the IO type that is most problematic for flash writes. VAST also uses this buffer to manage how data is written to the third key technology QLC flash. The key for VAST is the large storage class memory buffer gives them the luxury of time to optimize how the system will interact with QLC.

3

VAST uses a series of methods to maximize the life expectancy of QLC technology. The Storage Class Memory buffer, today leveraging 3D XPoint to buffer writes enables VAST to perform full-stripe writes to the QLC buffer. This minimizes the write amplification typical in all-flash systems and which is especially problematic for QLC media. The data is also similarity compressed prior to being written to QLC, and because the data is stored natively in the storage class memory buffer, they can maximize data reduction efficiency. This is something that VAST calls similarity-based reduction. Since the system is also file system based it can monitor the file types most associated with temporary files and ensure that it does not intermingle that data with more long-term data.





VAST's data protection capabilities are also impressive. First, it leverages an enhanced data protection algorithm that takes advantage of the solution's wide striping to provide maximum protection with minimal capacity overhead. Typical raid protection groups are 150+4 which VAST claims is 3 times faster than HDD erasure coded rebuilds with only 2.7% capacity overhead. The 500+8 grouping is 2 times faster than HDD erasure coded rebuilds with only a 2% capacity overhead.

In addition to VAST's media protection capability the system also provides continuous snapshots capabilities. The system uses a redirect on write technique to provide high performance, with nearly unlimited snapshots. Snapshots are also granular to the file level because of its file-system foundation. Snapshots are writable and cloneable and volumes can be rolled back to any protection interval.



The 150+4 grouping is **2 times** faster than HDD erasure coded rebuilds with under **3% capacity overhead.**

STORAGESWISS TAKE

The VAST solution is technology rich and is much more than a solution that is QLC ready. VAST has essentially re-thought every aspect of what has become typical in storage system design and leveraged the latest innovations to deliver a system that promises an unprecedented price and performance ratio. The solution is one of the first to compensate for the reality that all-flash gives data centers more than enough storage performance and delivers what they really need, storage simplicity and faster time to insights.



Storage Switzerland is the leading storage analyst firm focused on the emerging storage categories of memory-based storage (Flash), Big Data, virtualization, and cloud computing. The firm is widely recognized for its blogs, white papers and videos on current approaches such as all-flash arrays, deduplication, SSD's, software-defined storage, backup appliances and storage networking. The name "Storage Switzerland" indicates a pledge to provide neutral analysis of the storage marketplace, rather than focusing on a single vendor approach.



Vast looks past the marginal gain and apply unconventional thinking to look at decades-old problems through an entirely different lens. Their mission is to bring an end to decades of complexity and application bottlenecks that have been caused by mechanical media and by the complex tiering of data across different types of storage systems. To achieve their goals, Vast reduces the problem in order to achieve exponential gains. The result: a dramatically simplified customer experience that marries with the ability to compute on vast reserves of data all in real time. Over three years, the VAST story has transformed from concept to reality. Since releasing V1 of VAST Data's Universal Storage concept in November of 2018, VAST has now established itself as a leader among the fastest growing IT infrastructure companies of all time.