Flash Storage Buying

A guide for IT Professionals

Contents

- + INTRODUCTION
- + PERFORMANCE
- + DENSITY & SCALE
- + RESILIENCY & AVAILABILITY
- + EFFICIENT MANAGEMENT & CONTROL
- + FUTUREPROOFING ARCHITECTURE
- + TOTAL COST OF OWNERSHIP





Introduction

With the introduction of enterprise-grade all-flash arrays, conventional thinking about storage has seen a serious shift in the mindsets of IT professionals. No longer is flash storage being seen as too expensive; it now enjoys 'must buy' status in evaluating storage purchase.

Industry analysts find that on average, all-flash storage pays for itself within six months, offering dramatic improvement in total cost of ownership (TCO) due to the fact that flash storage reduces operating costs associated with administration, power, space, cooling, and maintenance.

Because of these TCO reductions, Wanstor has seen many of its customers giving serious consideration to moving their storage estates to flash. At Wanstor we believe that flash storage may provide a solution for IT teams struggling to meet storage SLAs with application support, accelerated data generation and increased user demand.

Some customers have told us how their old tape based storage systems are struggling under existing workloads. With GDPR on the horizon, IT teams need to plan for both future storage requirements and changes in the regulatory landscape. There has never been a better time to consider switching to all-flash solutions. We have developed this guide to assist IT professionals in understanding what's important in flash storage purchasing and the options available to them.

Key areas covered in this guide for flash storage include performance, capacity and scalability, resilience and availability, future support needs and total cost of ownership.

By using these factors as criteria for making a purchase, IT professionals should be able to undertake an objective evaluation of their storage requirements.

3

Let's get started.

Performance

The majority of available all-flash storage systems make use of solid-state drives (SSDs) available from various vendors. SSDs are generally seen by most IT professionals as delivering superior performance versus hard disk drives (HDDs), in terms of three common storage performance metrics - I/O Operations per Second (IOPS), Throughput measured in MB/sec, and Latency, measured in milliseconds or microseconds.

When considering all-flash storage options, think about the I/O characteristics of the workload to be supported. All-flash storage systems that offer quality of service (QoS) controls can provide significant advantage in situations where programmes of work are taking place to consolidate multiple applications on a single storage system, preventing any one application from consuming too many resources, thus allowing the IT team to prioritise importance.

Transaction-oriented applications such as databases generate small random read / writes; this type of storage performance is measured in IOPS, and is most commonly reported for all-flash storage.

For meaningful IOPS comparisons it is important to know operation size and the mix of read / write operations. This metric is where SSDs and all-flash arrays differentiate most from traditional HDD-based storage. Single SSDs can deliver IOPS measured in tens of thousands, whereas a single HDD can deliver only hundreds.



figure 1 : Projection of Revenue generated (GBP) worldwide by enterprise flash storage from 2017 to 2022

PERFORMANCE

4

Most vendors publish IOPS ratings for their all-flash arrays, so it is always worth comparing different IOPS rates to see a true reflection of performance at a baseline level, and evaluate whether the technology warrants further investigation.

Latency describes time taken for a storage system to satisfy an I/O. It is an important metric for time-critical applications such as online transaction processing, betting, and financial trading, all of which are sensitive to latency.

Users react to differences in latency as proven by Amazon. com, who claim their tests showing a 100ms increase in latency corresponding to a 1% drop in sales revenue.

Although many all-flash storage systems support latencies as low as 1ms, applications that require extreme performance may require consistent, sustained response times that are measured in hundredths of microseconds rather than milliseconds. These applications do not tolerate latencies that fluctuate due to backend storage services, or that increase rapidly as IOPS load increases.

Throughput applications such as data warehouses and video rely on sequential access to data in large blocks, resulting in workloads that are dramatically different from transaction-oriented applications. Throughput measures the amount of data that can be moved in or out of a storage system and is typically reported in MB/ sec or GB/sec. Not all vendors report throughput numbers. If IT Managers know they have a throughput-oriented application, they should make sure to fully evaluate throughput performance.

As performance and density of all-flash storage systems improves, likelihood of network bottlenecks also increases, making network performance an important consideration.

Wanstor has seen many customers with I/O performance bottlenecks already transplanted to the network from storage layers

Upgrades should not be undertaken in isolation, but should include other areas of infrastructure to see where issues or challenges may occur after upgrading. We have observed wide performance variations across vendor systems and even variation within single systems, depending on factors such as differences in read and write performance.

Generally, Wanstor's advice for IT Managers is that they seek solutions which demonstrate consistent, scalable IOPS performance at latencies under 1m/s based on third-party benchmarks that simulate real world workloads, such as SPC-1 and TPC-E. Additionally, they should be wary of vendor claims regarding maximum IOPS without detailed proof.

5

PERFORMANCE

Asking the following questions will help in comparing differing vendor IOPS claims:

- What was the I/O size used to run tests?
- Were I/Os random or sequential?
- How many IOPS are supported at 1m/s latency?
- How does this compare with the I/O size for my applications?
- What was the read / write mix?
- What QoS features are available to make sure that IOPS are available when you need them?

Can your network keep up with all-flash storage density?

Few SSDs can saturate typical 8Gb and 16Gb Fibre Channel SANs or 10 Gigabit Ethernet networks. When all-flash storage system density increases, **network bandwidth requirements for associated storage systems should increase in line.**

Unless you plan and deploy storage purchases ahead of time, network bottlenecks will result. Look for storage that supports multiple network connections and current highspeed network options, including **32Gb Fibre Channel** and **40 Gigabit Ethernet**.

Density & Scale

One reason we see IT teams updating to all-flash storage is increased density and scale, allowing:

Greater flexibility : With the IT environment evolving rapidly, teams need to make storage choices that offer greatest flexibility for future movement in supporting the business.

Increased density: Replacing older hardware with all-flash systems may release valuable data centre space and decrease power and cooling costs.

Infrastructure standardisation : By eliminating specialised storage infrastructure supporting different applications and standardising hardware selections, teams may create a simpler, more flexible data centre that adapts to change more readily. Storage density of SSDs is growing faster than HDDs. In 2016, 15TB SSDs were released, exceeding capacity in the densest HDDs available at that time by more than 50%.

Practically, this means that a small number of SSD storage racks replace multiple HDD storage racks. This means recovered space and power and cooling reductions.

Not all storage systems can accommodate newer high-density SSDs; if maximizing capacity is your goal, this should be part of your decision criteria.

When comparing all-flash storage, note that many vendors report two capacity metrics: raw capacity and effective capacity. The effective capacity measures data the vendor expects storage systems to hold after deduplication and compression are applied.

Most all-flash arrays today provide inline deduplication and compression, reducing total storage to be purchased for storage. The effectiveness of storage efficiency technologies differ between vendors, so be critical of vendor claims when vetting options.

Deduplication and compression stats are not always key when evaluating all-flash systems. For applications demanding lowest possible latency, eliminating inline storage efficiency technologies reduces the data path, lowering latency and making I/O delivery more predictable.

There are generally two approaches you can take when scaling all-flash storage: scale-up or scale-out

With the familiar scale-up method, administrators add SSD storage to a dual-controller storage system until they reach capacity or the limit of performance. Thereafter, options include replacing existing storage systems with more powerful hardware or adding separately managed systems. When scaling-out, storage controllers are clustered together such that additional controllers can be added to scale performance and capacity beyond the limits of a conventional dual controller system.

This offers a single storage pool and simplified management versus multiple independent systems. Scale-out architectures eliminate the complexity and disruption from scaling capacity. This is the approach that we recommend.

Resiliency and Availability

Although industry data indicates SSD technology as inherently more reliable than HDDs, they are not completely fail-proof; standard forms of data protection should apply. Most all-flash storage systems incorporate a redundant array of independent disks (RAID). Given their superior performance, should an SSD within such an array fail then rebuilds happen far faster, limiting exposure to a second failure.

In many situations, flash storage provides enough of a performance cushion so that users and customers are unlikely to see a change in the data services they consume during a rebuild.

Reliable, all-flash storage systems remain subject to both hardware and software failures, leading to unplanned downtime. Wanstor recommends selecting proven SSD vendors with robust architecture, mature processes, historical reliability and well established support in place.

7

Improved Reliability

Concerns about flash media's durability or 'wear-out' still exist, even though the data attests to its reliability. Vendors are often so confident in system robustness that these ship with extended warranties; additional offers may entice you into long term service contracts. Factor these 'offers' and discounts into TCO calculations when making a direct comparison amongst vendors. Potential savings from some of these are offset by higher annual service charges.

As data centre technology progresses, planned downtime for maintenance and upgrades becomes almost a thing of the past. Seek out architecture that supports non-disruptive operations for all common maintenance activities when moving to flash storage solutions. Backup and disaster recovery (DR) is standard when protecting against user errors, software bugs, widespread power outages and other natural or man-made disasters.

Modern storage architecture typically integrates snapshot and replication functionality that makes backup, recovery, and DR functions faster and more efficient. A mature all-flash storage system should include such features, along with asynchronous and synchronous replication, application-level integration and support for an ecosystem of data protection partners.

Efficient Management & Control

At Wanstor we understand that IT infrastructure is a means to an end, not an end in itself. The continuous monitoring, management and maintenance required by traditional storage systems keeps many IT teams from focusing attention on applications and services long enough to move business forward. We suggest flash storage options which:

- Offer fast and efficient data management capabilities including snapshot copies, replication, and cloning
- Provide intuitive user interfacing with role-based access control
- Reduce or eliminate routine IT operational tasks
- Enable maintenance and upgrades without planning for downtime
- Integrate with existing management and monitoring frameworks
- Provide REST APIs, plug-ins and other integrations, simplifying automation and storage in existing data centre environments

EFFICIENT MANAGEMENT & CONTROL

8

Futureproofing Architecture

The storage solution you choose will have to adapt to future needs. Storage purchased today should support the changing needs of your business, existing applications, new cloud applications and upgraded operating systems. It must be a good fit for current application environments and integrate with enterprise applications, such as those from Microsoft, VMware, and SAP.

This will facilitate provisioning and application-consistent data protection, and could include support for faster network technologies and other infrastructure upgrades.

Other considerations must include support for different storage protocols, or the replication of data from flash storage to other media, either on-premises or in the cloud.

Cloud is, of course, another critical area of application performance; next-generation cloud applications and services are typically optimised to run in the public cloud or over a private cloud or enterprise cloud infrastructure.

These environments also require all-flash capabilities described in previous sections covering performance, management and resiliency. In addition, they benefit from scale-out, easy automation, full REST APIs and Quality of Service.

FUTURE-PROOFING ARCHITECTURE



If a business develops or is planning on developing next-generation applications, these must run both on-premises and in-cloud. Storage solutions offering cloud integration can simplify the migration of applications. Consider investment into all-flash systems that are capable meeting of the following criteria:

- Data tiering between on-premises and cloud storage
- Movement of workloads to and from the cloud
- Backup of on-premises storage to the cloud
- Disaster Recovery within the cloud

Total Cost of Ownership

For most IT departments, cost is a mitigating factor in most purchasing decisions. Various studies have found all-flash storage significantly lowers TCO relative to traditional storage systems.

To validate this, Wanstor can help in comparing different solutions. Our flash storage comparison workshops take into account vendors, performance, operational costs and initial capital outlay.

IT Managers can see how swiftly a flash storage solution will repay the business compared to existing disk-based or hybrid flash systems when justifying a purchase. Additionally, effective capacity should be factored into the decision-making process as well. At a minimum, calculations should include inline efficiencies available with an all-flash system, such as data compression and deduplication, as well as snapshot copies and clones (if your business uses them). Management costs for provisioning, protecting, and recovering application data are also often overlooked when evaluating primary storage.

The capabilities and costs for managing application-consistent data copies can vary widely across all-flash offerings, and should be included when calculating TCO. Data centre operational costs in terms of power, cooling, and real estate are also factors.

In user surveys, flash storage solutions save between 25% to 40% in power and cooling when compared with diskbased systems

Database, application and infrastructure specialists can all reclaim time that was previously spent architecting and managing storage performance; this in turn will help in consolidating workloads onto an all-flash platform as well lead to software savings.

The final consideration should be operational management costs. When storage performance no longer need be treated as a scarce resource, IT personnel are released from the bonds of hardware across the entire business.

TOTAL COST OF OWNERSHIP

10

Final Thoughts

When evaluating flash storage solutions, businesses should ensure they undertake fair comparison of all technology available compared to other storage solutions.

At Wanstor we provide a range of flash storage services, from technology evaluation workshops to best practice advice on procurement and in-life management.

The advice in this document along with our expertise in flash storage means many businesses already enjoy better storage performance, improved application management and reduced IT operating costs.

All our flash storage solutions are built with your business in mind. The solutions we deploy and manage for customers enable their businesses today and allow them to stay ahead of the competition whilst remaining compliant with changing data management regulations.

We hope you've enjoyed reading this buyers guide, and look forward to discussing your future flash storage procurements soon.

Wanstor | 124-126 Borough High Street | London | SE11LB | info@wanstor.com © Wanstor. All Rights Reserved.