

Everything I need to know to evaluate **flash**

Flash Buyer's Guide for Enterprise Storage

Foreword by Ben Woo, Managing Director, Neuralytix



Table of Contents

1. [Foreword by Ben Woo, Neuraletics](#)
2. [Introduction to the Flash Buyer's Guide](#)
3. [4 Things to Know About Flash Memory](#)
5. [The 3 Classes of All-Flash Products](#)
5. [What is Your Application Footprint?](#)
7. [How Does the Product Protect the Flash?](#)
10. [What Data Reduction Does It Offer?](#)
11. [What Serviceability Does It Offer?](#)
12. [What Application Integration Does It Offer?](#)
13. [How Mature is the Product?](#)

[4 Things to Know About Flash Memory](#)

Before we dive into all-flash storage, there are some fundamental attributes of flash memory with which you should be familiar...

[The 3 Classes of All-Flash Products](#)

SLC (single-level cell)... eMLC and MLC (enhanced multi-level cell)...
TLC (triple-level cell)...

[What is Your Application Footprint?](#)

The purpose of this section is to determine the applicability of your application workload for all-flash storage.

Foreword: Neuralytix

By Ben Woo, Managing Director, Neuralytix

All technologies evolve. From time to time, technologies will move from being evolutionary to revolutionary. This is the case with flash-based storage.

With the right solution, and the right deployment, flash storage can improve competitive advantage significantly, and drive down capital and operational costs of IT. The speed of flash yields application improvements. But enterprises can also benefit from flash performance by accelerating time to insight or time to value. In both of these cases, enterprises who are late to deploy flash will find themselves at a distinct disadvantage.

However, it is not too late to integrate flash into your enterprise infrastructure.

This guide provides the questions that will serve to help prospective all-flash storage systems buyers to compare and contrast the various solutions presented to them. It will also help to understand the various nuances that distinguishes flash-based storage from traditional hard disk drive based storage. The questions asked in this guide can also be used by prospective buyers to formulate a RFP or RFI for all-flash storage.

Hybrid storage solutions have been excluded in this guide as most hybrid storage systems are traditional hard disk drive based storage systems with the option to integrate solid state flash memory that is packaged in a disk drive form factor. This approach certainly provides the opportunity to add a faster tier of disk, but does not necessarily have distinguishing features or benefits like an all-flash array does.

Many vendors have jumped on the flash bandwagon giving end-users a wide variety of choice. While choice can be a positive aspect of an active market, it can also reflect fragmentation and cause confusion to end-users looking to benefit from new technologies.

For example, flash memory, the core of flash storage arrays is available in multiple flavors ranging from consumer grade to enterprise grade flash. Flash storage is also available in multiple form factors, with varying degrees of performance benefits. Some flash storage solutions will benefit a singular application or server, while others will provide benefits across a plethora of servers with a diverse range of workload and performance needs.

All flash storage systems represent a new approach to traditional magnetic disk based or hybrid media storage systems. But not all flash storage systems are designed the same way or yield the same benefits. This guide provides an end-user with a collection of key criteria to be used in evaluating all flash storage systems solutions.



Ben Woo
Managing Director, Neuralytix



All-Flash Storage Buyer's Guide (2013 – 2014)

All-flash storage can be a competitive advantage for your entire company – it has the potential to unlock a new level of productivity for your organization's employees by reducing the amount of time they wait for their databases, virtual desktops, and other applications. Flash's speed, space efficiency and reliability make it a highly attractive alternative to mechanical disk as a storage medium for enterprise applications.

In the last few months, a whole series of vendors have announced new products that qualify as "all-flash storage". The churn caused by these new entrants makes it complicated to make good decisions about storage purchases. Even if you haven't deployed flash in your infrastructure yet, chances are in the next year it's likely to happen.

So how do you make the right decision? With suppliers ranging from established incumbents to young new startups, and a variety of new technologies with which you may or may not be familiar, data in the market is imperfect and can result in a poor selection. There are also different product classes, each with a unique set of strengths and capabilities. Finally, flash memory itself introduces a different set of potential failure points – and technologies to address them – that are implemented in different ways but often are described under a common term. That can lead to even more mis-information for the potential buyer.

The purpose of this guide is to provide a basis for evaluating and selecting all-flash storage for enterprise class environments. It covers all product classes and where they can be best applied, as well as what the key elements are for each to avoid potential pitfalls in the selection process. We'll show you where you can be easily misled or confused by suppliers, and provide you with frameworks and tools you can use to make the best possible decision based on your environment and your applications.

This buyer's guide covers:

- The different classes of all-flash storage
- What you need to know about flash memory
- The buying criteria for all-flash storage based on your application

The number of IT professionals with hands-on experience with flash storage is increasing quickly, and with it comes a wealth of new – and sometimes confusing - information about the performance, reliability and general experience with this new class of products. We hope this buyer's guide offers some insights that will help you make the best decision for your organization.

Neuralytics Speaks on Flash This Year

"With the allure of the speed and efficiency that flash provides, users are bound to evaluate the onslaught of all-flash products introduced this year. By the same token, many of these products are neither trial nor time-tested. There will be much to learn for the storage buyer."

- Ben Woo,
Managing
Director

Four Things to Know about Flash Memory

Before we dive into all-flash storage, there are some fundamental attributes of flash memory with which you should be familiar. We will assume you are familiar with the beneficial attributes of flash – its speed, power and space efficiency and reliability are generally known. This section speaks explicitly to some of its weaknesses and how it's different than mechanical disk.

1. Flash Cells Wear Out

Flash memory is a non-volatile semiconductor product that was originally designed for digital cameras. These cells were not intended to be written to in the intense and frequent fashion that most enterprise applications do this. One of the challenges with this is that flash cells wear out a little every time they are erased/programmed. The expected life of a flash device is expressed in program/erase (PE) cycles. A good analogy for understanding flash wear is to think of a piece of paper. If you were to erase a piece of paper in the same spot 1,000 times, surely you would tear through the paper. This is the same way a flash cell works.

You'll see in the following table that the lower the cost per bit, typically the lower the number of program/erase cycles a flash device is warranted for.



Figure A: How a Flash Cell Wears Out; Same as Erasing on Paper

2. Flash Comes in Different Flavors

	SLC	eMLC	MLC	TLC
Program/Erase Cycles	100K's	10K's	1K's	100's
Performance	Fastest	Middle	Middle	Middle
Cost per Bit	Highest	Middle	Low	Lowest
Bit Density / Bits per Cell	Lowest / 1	Middle / 2	Middle / 2	Highest / 3
Bit Error Sensitivity	Low	Low	Low	High

While the fundamental blueprint of a flash cell is basically identical, how they are manufactured and screened for performance and other characteristics is very different and drastically affects the behavior, cost, performance and reliability of the end system that is built on it.

- SLC (single-level cell): this is the most expensive form of flash memory because you can only store one bit in every cell. It is also the highest performing, and approaches DRAM in terms of the speed of transactions it can support. It also has the highest rating for endurance.
- eMLC and MLC (enhanced multi-level cell): these two styles of flash are almost identical except for endurance. MLC increases the bit density by making it possible to store two bits per cell

instead of just one, and is the type of cell used in the vast majority of consumer applications today, making it the highest volume flash device today. eMLC is a specially-screened device derived from MLC that simply increases the P/E cycles at the expense of the product's yield, which increases its cost.

- TLC (triple-level cell): the most recent innovation in the flash market, TLC allows for 3 bits to be stored in each cell. However, the architecture has two material weaknesses; first, it is far more susceptible to bit errors due to the narrow band in which the bit's value can be read. Second, it wears out very quickly, supporting merely hundreds of write cycles.

3. Flash Fails Differently than Mechanical Disk

If you extend the notion of how flash wears to bit error rates, you will find that flash memory fails in a very different way than a mechanical disk;

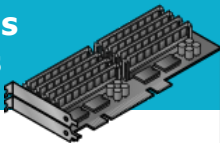
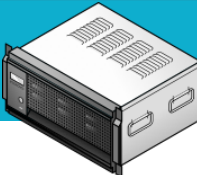

- Mechanical disks experience bit errors with high adjacency – think of a scratch on an LP
- Flash experiences bit errors based on variables such as wear and data retention – the media gets worn over time

Further, the unrecoverable bit error rate for hard disk drives remains relatively constant over the number of writes to the drive, while with flash, the error rate increases exponentially with the number of TBs written. For some period of writes, the flash solution will have a much lower bit rate, then it will rise dramatically to meet that of the HDD. This fundamentally different slope of error means that unique error correction schemes need to be applied for flash – schemes that aren't appropriate for an HDD.

4. Flash Writes are Expensive, But Reads are Free

Unlike a mechanical disk, flash reads and writes are highly asymmetric; it takes much more time to erase and write a flash cell, than it takes to read. The consequence here is that when architecting a system around flash, this asymmetry is key to ensuring predictable high performance as you want to minimize the number of writes you are performing, and not worry about reads at all.

The Three Classes of All-Flash Products

PCIe Cards in Servers 	Flash Appliances 	Flash Arrays 
<p>Purpose: Acceleration of single applications</p> <p>Key attributes:</p> <ul style="list-style-type: none">• The absolute fastest flash-based IO acceleration possible• No software for high availability: must mirror arrays, doubling cost of flash capacity• Relative low capacity (<2TB) results in IOs hitting the shared storage incurring both line and disk latency• Typically requires heavy optimization of the application• Most expensive option on a cost-per-GB basis <p>• 1s of TBs</p>	<p>Purpose: Acceleration of specific applications</p> <p>Key attributes:</p> <ul style="list-style-type: none">• Combines very fast performance with high capacity• Typically limited or non-existent enterprise storage features (NDU, HA, snapshots, etc).• Dense / compact per U <p>• 10s of TBs</p>	<p>Purpose: General-purpose Tier 1 storage</p> <p>Key attributes:</p> <ul style="list-style-type: none">• Highest capacity, enterprise-class features• Not the absolute fastest, incurs line latency to the storage array• High availability, non-disruptive upgrades, snapshots, other enterprise features included in software stack.• Scale out to the largest capacity of these categories <p>• 100s of TBs</p>

Before You Begin:

Know Your Application Footprint

The purpose of this section is to determine the applicability of your application workload for all-flash storage. An understanding of your workload will help you to determine what an investment in all-flash storage will do to improve your results.

How sensitive to latency are your applications?

If your applications can tolerate tens of milliseconds of latency, you may not need a flash solution at all. If your applications need predictable, consistently low latency, then all-flash solutions may work well for you. But be warned: many all-flash architectures are not designed for predictable latency. If the all-flash architecture uses SSDs, be aware that SSDs suffer from occasional “performance brown-outs” where the SSD will go to sleep temporarily while it executes internal procedures. There is also the asymmetric nature of reads and writes that need to be addressed to produce predictable latency.

What does your access pattern look like?

Do you understand what your read / write ratio is? Virtual desktop environments tend to be very write-centric, while databases are more read-centric. Understanding this ratio will help you to

make a better storage decision around all flash. In general, because flash reads are so fast, the more read-centric an application is, the greater the latency benefit you will see. This is not to say that write-intensive applications will not see a benefit, it will just be less dramatic.

What does your IOPS load look like?

How many IOPS are your applications really using? While you might hear storage suppliers claiming they can do a million IOPS – that might be far from what you need. If you are in the tens of thousands of IO, on this factor, you probably don't need to upgrade to flash memory for your application; mechanical disk can handle these smaller loads fairly well.

What is your typical block size?

Most benchmarks assume a 4KB block size. The primary reason for this selection is because modern hard disk drives are broken up into 4096 byte sectors to solve a number of hard disk drive problems – it allowed for more efficient storage of large files, at the cost of less efficient storage for small files, and enabled the integration of stronger error correction codes that were needed at the higher storage capacities that hard disk drives were achieving. In other words – it didn't really have anything to do with what applications were doing. It also has nothing to do with flash architecture.

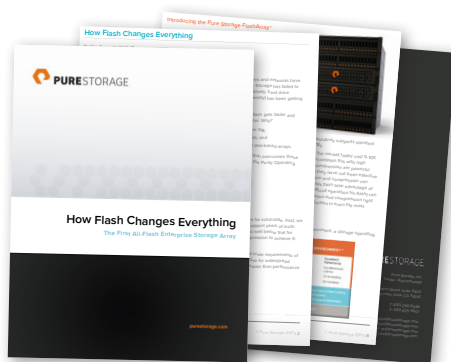
Interestingly, many all-flash arrays are using a 4KB block size.

Why? It may be because the designers believe that users have optimized their architectures to the 4KB block since its introduction on mechanical disks a few years ago. What you should know is that benchmarks that use a 4KB block size are irrelevant to your workload unless you've already re-written your application to optimize for this size.

Perhaps the most important reason you need to understand your block size is due to block alignment issues.

Effectively, if you have large block sizes in your flash array, those blocks can fail to align well with the applications blocks, resulting in a reduction in performance.

The reality is that block sizes are highly variable based on the applications you are running. At Pure Storage, all-flash storage arrays report the block sizes of the datasets running on them; the average block size is 64KB!



Download How Flash Changes Everything.

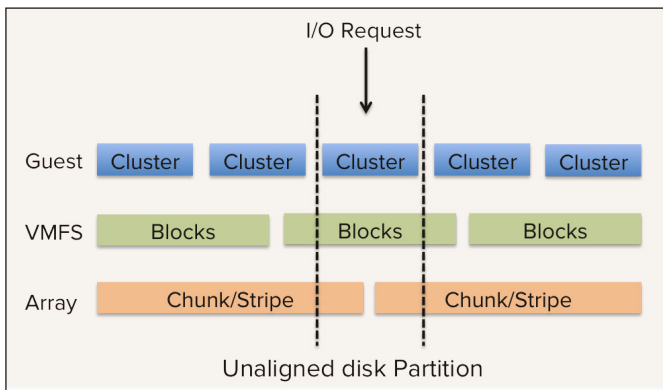


Figure B: How block alignment affects performance

What is the locality of your access?

Access locality refers to the relationship between your IOs. Are they sequential and predictable? This is more common in environments where you have, for instance, a large stream of video being read in order from the storage. The opposite end of the spectrum is a highly random and unpredictable access, where each IO in the stream may be requesting or writing unrelated data. Random IO streams are commonplace in any virtualized environment due to multiple VMs IO streams being interleaved to the shared storage array.


The more random and unpredictable the IO stream, the more of a performance advantage a flash storage solution is going to have, because flash is inherently random in its ability to access data.

How to Qualify a Workload for Flash

Now that you know what your workload looks like, it's time to map that workload to see which storage media is best suited to your objectives.

To the right is a figure that summarizes the IO footprint variables to help you understand which types of workloads will do best on flash.

Figure C: Classifying if a Workload is Well-Suited to Flash Memory



How is the Data Protected?

Flash memory, due to its performance, is most frequently used in tier 1 applications that require both a high level of data protection and performance. While the different methods of flash deployment mentioned earlier in this document are valid, their ability to provide data protection varies wildly. Usually, this requires a material amount of software in the system to provide the fault tolerance that these applications need.

What type of RAID protection is being used?

RAID architectures were invented for mechanical drives, not flash. So direct application of RAID levels to flash has to be understood at a level slightly deeper than a simple handle like RAID 5 or RAID 6.

Server flash cards will really only be able to provide data protection via mirroring due to their insertion directly into the server. This means that if a flash card fails, the data on it has to have been saved on a separate card and the IO traffic needs to be re-routed to the other card.

In effect the biggest consequence here is the cost – this architecture doubles the cost per bit. Shared storage architectures (NAS/SAN) will use various types of RAID including mirroring, RAID 5, and RAID 6. How a RAID type is determined includes many factors including rebuild time, application read/write mixture, performance requirements, and required usable capacity. After various complex calculations including frontend and backend IO requirements due to the different write penalties, risk discussions regarding rebuild times, and capacity analysis, a RAID structure can be defined and implemented. As mechanical drives become more and more dense, rebuild times monopolize the RAID decision. This is because the window during which the rebuild is occurring is a point of failure should another drive fail. That said RAID 5 can be very risky, so in most cases mirroring or RAID 6 are the only options.

With flash storage, due to the incredible performance of the storage medium, the rebuild duration should be materially shorter – and have only a small impact on the performance of the array. Flash's natural architecture lends itself to the uniform distribution of parity bits across a number of drives. For any architecture greater than 10TBs, it is critical to have support for dual-parity as the propensity for lost bits increases as the array scales in size.

Another significant difference between mechanical disk RAID and what flash RAID can do is the way a RAID rebuild is executed. With mechanical disk/traditional RAID, the architecture is a many-to-one rebuild (RAID set to a single hot spare). With flash, it is possible to design RAID to execute a rebuild in a many-to-many configuration, using the free capacity of all drives being rebuilt to.

A major issue with mechanical disk arrays is how long it takes them to rebuild when a drive failure happens. This is because the window during which the rebuild is occurring is a point of failure should another drive fail. With flash storage, due to the incredible performance of the storage medium, the rebuild duration should be materially shorter – and have only a small impact on the performance of the array.

In summary, you need to know how many bits of parity exist in the RAID architecture, how many drives or cards the storage uses for a RAID group, how many drives can fail concurrently without losing data, and how long it takes to rebuild a drive under specific circumstances (50% full, for instance).

How much overhead does the RAID require?

Implementing the parity bits within a RAID stripe consumes space on the array. An understanding of how much additional space the parity bits consume is important to understanding what the fully-protected cost per bit for your all-flash storage will be. This should be expressed as a percentage of the physical storage consumed.

Does the architecture use hot spares?

Hot spares in a flash storage system are powered on but not populated with data. Additionally, hot spares are subject to a window during which the original device has failed, and the hot spare is being populated with data, during which a subsequent drive failure can cause data loss. So while they improve the resilience of the architecture, they are not fail safe.

The presence of hot spares introduces additional cost to the system based on the protected capacity, so like RAID overhead, it's important to understand what percent of the usable capacity

the hot spares cover and the cost for that coverage.

How many drives can fail at once?

If the all-flash storage system is a shared storage solution and contains either SSDs or an array of flash cards, an understanding of how many failed SSDs/cards you can lose at once before a data loss will happen should be a key selection criteria. Dual-parity systems should be able to support at least two drive failures within an array.

And what happens when a drive or card fails?

The supplier should be able to explain the rebuild and replacement process. How long does it take for the drive or card's data to be rebuilt from parity? What impact does the rebuild have on the performance of the array during the rebuild process?

Some suppliers will require you to remove the equipment from the rack in order to replace a drive. This often is accompanied by the need to turn off the device and interrupt service.

How do you protect data as it is received by the array?

It is also helpful to understand what happens if a cache fails in the system – when an IO request is placed to the storage, where is that data stored until the array responds and how is it protected from loss? Is it stored in a non-volatile location or just in RAM?

What level of serviceability does the architecture provide?

When a drive or component fails, how difficult is it to replace the component? Some flash storage solutions require you to pull the equipment from the rack and open the chassis in order to replace a failed device. This is true for PCIe cards, whether placed within the server or in a dedicated storage appliance.

An important factor to keep in mind is how fast SSD generations change – namely every 12-18 months, you have an entirely new group of SSDs in the market. If this is an SSD based architecture, how do you replace an SSD that's been made obsolete? What if that component can no longer be sourced?

Does the architecture support snapshots and replication?

A key enterprise feature of shared Tier 1 storage is the ability to take snapshots and replicate the data to other sites for disaster recovery. On flash memory, you can expect much higher performance for these processes, especially if the system uses deduplication / metadata, as creating snapshots can simply be producing an identical set of metadata pointers, as opposed to actually writing the data anew.

If a flash storage solution does not include replication, there are some third party options that can be considered that allow array-to-array replication but are agnostic to the array supplier. While these add licensing costs, they do solve the problem if it's a requirement in your environment.

What Methods of Data Reduction are Offered?

Data Reduction Types: Deduplication, Thin Provisioning, Compression, Pattern Removal, or All

You will find a number of different methods of data reduction in the market around flash storage – these should be not treated as check box items as they have a material effect on the performance, efficiency and cost-per-bit of your storage.

Different vendors have some but not other data reduction technologies in their systems. It's important to understand the origins of the technology – if it wasn't intentionally built for flash memory, it likely has a material impact on the system's performance. Some all-flash appliances offer no data reduction at all. Below are some factors to be aware of when selecting a flash solution.

Inline vs. Post-Process Deduplication

Some vendors will reduce data as it is coming into the array, others will first land data to disk and/or flash before processing for data reduction. In the legacy disk world, post-process was the norm because data reduction took a long time to execute with available disk and CPU performance. In flash, inline deduplication becomes possible due to the performance of flash and new architectures designed for it.

It is important to understand that inline affords two key advantages:

1. Cost reduction via increasing the effective capacity of the flash, and
2. Write avoidance by removing write IOs that would have otherwise hit the flash

Reduction: "Chunk" Size, Granularity, and Alignment

Data reduction technologies vary in the granularity with which they analyze data. Granularity is a bit of a trade-off: the smaller chunks you use to analyze data the more redundancy and benefit you'll find, but the smaller the chunks the more metadata the process creates that must be managed. Granularity also has a direct tie to alignment...anyone who has ever dealt with mis-aligned VMs in traditional storage knows exactly what I'm talking about...if all the layers (storage, hypervisor, VM, FS, application) don't align on the right geometry, finding duplicates is all the harder. Most data reduction and thin provisioning technologies today operate on the 4K or larger chunk size.

Performance Impact

In traditional disk storage dedupe is really slow and just can't be run in any performance environment. In fact, the same challenges exist in poorly-designed or retrofit data reduction implementations in the flash world as well: enabling data reduction can come with a 50%+ performance impact. If a vendor is telling you about their data reduction, the performance impact that comes along with it should be your very first question.

Global vs. Local

Many of the existing implementations of data reduction are confined to a portion of the storage array to keep their performance and metadata manageable, such as only within a volume, LUN, FS, shelf, etc. The downside is that the more you divide data reduction into pools or zones, the

more you duplicate data between these pools and reduce overall effectiveness.

Purpose-built vs. Retrofit

To enable data reduction technologies to work at flash speed and global scale, it is necessary to design from the ground up to be optimized for data reduction. As vendors have realized that data reduction is a “must have” feature to make MLC flash arrays work, vendors may OEM or acquire 3rd-party technologies and retrofit them to their storage arrays, often as a post-process step.

Environmental Factors

Keep in mind that if you use data reduction at the host or application-level, you may see a reduction in the benefits offered by your all-flash storage. If you choose to use the application side reduction, you are using host CPU cycles to perform these tasks, and you cannot get the benefit of a higher level of dedupe across multiple hosts.

The flash storage will need to execute the deduplication regardless because it needs to minimize flash writes.

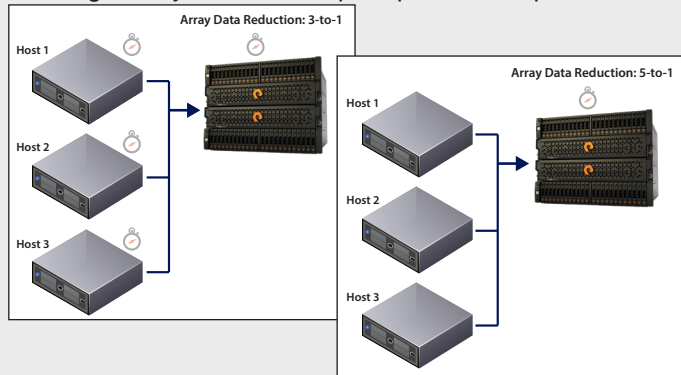


Figure D: Release CPU cycles from the host and let the storage array execute the dedupe

What Level of Availability is Offered?

Is it active/active, active/passive, or something else? The level of availability depends largely on the architecture of the storage system. If it is built to be a typical storage array, the consequences of a controller failure and how it is remedied is something to know before you buy the product.

Active/Active IO Handling

Mission critical applications typically cannot tolerate outages, so an active/active configuration is the best solution in these environments. These applications are typically the most performance-sensitive so they are often the primary motive for an all-flash storage evaluation.

In an active/active configuration, the protocol and bandwidth of the ports used to connect the two controllers enables the transfer of all IO traffic from a failed controller or controller component to the other controller. The speed of the failover will determine how disruptive the failover is to IO that are in flight for the applications running on it.

How Disruptive is an Upgrade?

Most all-flash storage solutions do not support non-disruptive upgrades today. Determine whether it is possible to replace a controller or a component without having to take the flash system offline. Also, can the vendor upgrade the software on the controller without disrupting the operation of the storage? How long does this take?

Is the Controller Stateless?

If the vendor combines controller, its cache and flash storage into a single enclosure, you do not have a “stateless controller” environment. A stateless controller allows for the case where the

controller fails and you need to replace it, but don't necessarily wish to remove all the storage associated with it. By separating the storage from the controller, either can be replaced without disrupting the operation of the flash storage array.

What Application Integration Exists Now?

Virtualization Integration: VAAI Support

VAAI is a VMware API that allows ESX hosts to offload storage-specific processes to the storage array, freeing up precious CPU cycles on the servers for other tasks. The types of tasks include creating and locking files at the volume level, cloning blocks within the same array, "zeroing out" disk regions, thin provisioning LUNs and reclaiming space on the array. Some vendors have full support, others only partial, but it is common to simply declare "VAAI support" without further explanation of which of the primitives are supported. As a buyer, you should not assume that all are supported when presented with this claim.

Virtualization Integration: vCenter Plugin

vCenter is the main user interface that virtualization administrators use for managing their virtual environments. And most virtualization admins aren't really that interested in learning how to manage storage from a new user interface. vCenter integration should allow an admin to:

- Create LUNs
- Change their size
- Map VMs to LUNs
- Set vendor-specific path policies
- Monitor backend LUN performance

These tasks should be easily performed from within the user interface without having to copy and enter WWN numbers or other mundane tasks to map storage to the virtualized environment.

Virtualization Integration: Microsoft ODX

MS Offloaded Data Transfer (ODX) has similar benefits to VMware VAAI (offload functions from the host CPU to the storage array). The key benefits of ODX are:

- Copy offload across physical servers and virtual machines
- High-performance data transfer rate through the storage network
- Low server CPU usage and network bandwidth consumption during the offload read/write operations
- Intelligent data movement options allow applications to optimize the offload read/write process. ODX is relevant for any environment using Microsoft's Windows 8 or Server 2012. Microsoft provides a testing process; be sure the array you've selected has passed this certification.

How Mature is the Product?

In the world of all-flash storage, there's a wide range when it comes to product maturity. Some products have been in the market for a few years, while many others aren't available at all – although they are advertised. Understanding whether you are the first customer ever to lay hands on the product or not and how ready it is for your specific environment will help you make a good choice.

Benchmarks and Certifications

Unfortunately at the moment, many of the benchmarks in the market are designed to measure the performance of mechanical performance disk, so relying on standard published results on SPC-1 won't give you an accurate assessment of the performance or efficiency of a flash solution.

Also, many configurations on the compute platform side are optimized for the speed of mechanical disk, which can cause a lot of IO to queue at the server unnecessarily. To get an accurate assessment, active evaluation of the flash in your environment is the key.

That said, some of the application-level benchmarks will provide a reasonably accurate assessment. For instance, for VDI production deployments, both VMware and Citrix have comprehensive benchmarks that evaluate storage based on the application's performance running many desktops with a typical workload and duty cycle.

These benchmarks can be found at:

[VMware View Storage Benchmarks](#)

[Citrix XenDesktop Storage Benchmarks](#)

It can also be helpful just to know that other vendors have validated the compatibility of the storage with their environment.

[Here you can find the Hardware Compatibility List \(HCL\) from VMware for storage](#)

GA Milestone?

The GA milestone can be easily misinterpreted as every supplier has a unique definition.

Understanding when they shipped GA, and what exactly was included in the release helps to clarify how robust the product is. Vendors should also be able to tell you how many units they have deployed, in what markets.

Conclusions

Like all storage media, flash is not a perfect medium. But, like magnetic rotating disks, and linear tape before it, flash represents the next generation of storage performance. Flash performance is orders-of-magnitude faster than magnetic rotating disks.

The integration of flash into an enterprise's infrastructure will yield improved performance. In determining the best flash storage system and flash storage systems vendor, an end-user must ask the right questions not only of the vendor, but about the application needs and the areas requiring additional performance.

The flash market is nascent. As this guide shows, there are few benchmarks and certifications that objectively compare and contrast flash solutions. Using the criteria laid out in this guide will help the end-user to instrument where and how flash can improve the performance of an enterprise's infrastructure; and this will ultimately lead to competitive advantage.

More Please! 


Have questions or need more information?
[Email us at info@purestorage.com](mailto:info@purestorage.com)
and we'll help you through the process of
selecting the best all-flash solution for
your environment and workloads.

Pure Storage, Inc.
Twitter: @purestorage

650 Castro Street, Suite #400
Mountain View, CA 94041

T: 650-290-6088
F: 650-625-9667

Sales: sales@purestorage.com
Support: support@purestorage.com
Media: pr@purestorage.com
General: info@purestorage.com



PURE STORAGE

PURE STORAGE

PURE STORAGE