

FLASH STORAGE BUYER'S GUIDE

A Comprehensive Guide to
Evaluating Flash Storage



FOREWORD

by Eric Burgener

*Research Director, Enterprise Storage
International Data Corporation*

Driven by a rapidly evolving business climate and the need to service a broad mix of both legacy and new applications, enterprise storage requirements have forever changed. In response to this, enterprises have been rapidly integrating flash into the persistent primary storage tier, and to date well over 70% of organizations are using flash in mission-critical production environments. The resulting All Flash Array (AFA) market is one of the fastest growing segments of the overall external enterprise storage market with a five-year compound annual growth rate of 29.2%. By 2019, AFA market revenues will grow to \$5.6B and this segment will account for 20.5% of end-user spending on external systems, up from 5.8% in 2014. In terms of primary storage revenues, the AFA market will account for 60-70% of spending in primary storage by 2019.

Although AFAs were often initially deployed to meet the performance requirements of a single, very demanding application like a transactional database or virtual desktop environment, increasing customer confidence in the technology is leading to a significant shift. After seeing what AFAs offer in production environments, it is common for customers to start migrating other applications onto these platforms, thus setting them on the journey to their own all flash, internal cloud. An increasing number of new customers are considering the capabilities these platforms offer for dense mixed workload consolidation right up front, and this has become a key purchase decision criteria within the last 12 months. This trend has caused customers to focus on not only performance, reliability and availability, but also other key factors that support mixed workload consolidation like scalability, rich data services, an ability to integrate well into pre-existing data center workflows, and cloud-based multi-tenant management. Ultimately, this trend will lead to AFAs displacing hard disk drive (HDD)-based systems and becoming the enterprise storage workhorses for primary applications within the next several years.

Enterprises that are still evaluating flash based around acquisition cost per GB relative to HDD options are doing themselves a disservice. When flash is deployed at scale, as it is for mixed workload consolidation, the secondary economic benefits of flash deployment result in a total cost of ownership (TCO) that even last year was easily 50% to 80% lower than that of a comparably configured HDD-based system. These secondary economic benefits include far fewer devices needed to meet performance requirements, much lower energy and floor space consumption, fewer servers needed to drive storage performance, and lower software licensing costs (due to needing fewer servers). An increased focus on storage efficiency technologies like in-line compression and deduplication, thin provisioning, and space efficient snapshots and clones just within the last year has now brought even the effective acquisition cost per GB of AFAs in-line with or below that of HDD-based systems.

Although flash performance is generally at least 10x that of conventional HDD-based storage, the IT infrastructure transformation it enables is only part of the story. The business transformation is the real story with flash technology – customers significantly increasing revenues, improving customer service, enabling new processes that were never before possible, appealing to new customers, and opening up new markets. The majority of end users with whom IDC speaks about flash deployment experiences have stories to tell in this area, and these have strong appeal not only to IT but also line of business and C-level executives. As legacy enterprise storage platforms come up for technology refresh, IDC is strongly recommending that enterprises evaluate all-flash options.

This guide is designed to assist prospective AFA buyers evaluate the different options. It suggests key features buyers should look for, making it easier to compare and contrast the diverse set of market offerings that are available from startup and established vendors alike. The benefits and tradeoffs of different implementations are discussed, providing an excellent background for enterprises that will help them identify the solutions that best meet their own requirements.

INTRODUCTION

All-flash storage is a next generation infrastructure technology that can provide a competitive advantage for your entire company. It has the potential to unlock a new level of employee productivity and to accelerate your business by reducing 10 fold the amount of time databases and applications spend waiting for storage to respond with data. The “tipping point” for flash in the IT industry took place in 2014, as the storage market pivoted from mechanical disk to focus on delivering new products that qualify as all-flash storage. As the cost of flash continues to decline, flash has become the de facto enterprise standard for primary storage and the foundation for any cloud – internal, SaaS, or public.

Organizations of any size can now build towards their own all-flash cloud and we see the journey in three stages. The first stage is all about accelerating specific applications and experiencing the reliability and performance benefits of flash. The second stage is about deploying flash for all applications in an internal cloud that is fully virtualized and automated. Finally, the third stage of this journey is about using flash for cloud-native applications and for insights from real time analytics and big data workloads. Organizations of any size can embark on this journey at any time and experience the transformational benefits that flash can deliver.

Even if you haven’t deployed flash in your infrastructure yet, chances are it is already in your plans. So how do you make the right decision? With vendors ranging from established incumbents to young startups, and a variety of new technologies with which you may or may not be familiar with, data in the market is imperfect and can result in a poor selection. Flash memory itself introduces a different set of potential failure points, requiring new sets of technologies and architectural considerations to address. With different product classes, each offering a unique set of strengths and capabilities, how does one align the technology to their goals?

The purpose of this guide is to provide a basis for evaluating and selecting all-flash storage for enterprise class environments ranging from small enterprises up to the largest multi-national organizations.

This guide 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 confused by vendors, and provide you with tools to make the best possible decision based on your environment and your applications. This buyer’s guide covers:

- What you need to know about flash memory
- The different classes of all-flash storage
- Buying criteria for all-flash storage

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.



4 THINGS TO KNOW ABOUT FLASH MEMORY

Before we dive into all-flash storage, let's cover some fundamental attributes of flash memory. The performance and efficiency benefits derived from flash storage are generally well-known but flash is not mechanical disk and it has to be treated differently to ensure reliability and affordability.

1. FLASH CAN WEAR OUT

Flash memory is a non-volatile semiconductor technology that was originally designed for digital cameras. Flash cells wear out a little every time they are erased/programmed, which makes reliability a challenge with I/O-intensive, enterprise applications. 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. There are ways to minimize and even set the PE limits like over provisioning and write optimization; both are covered in greater depth in a bit.

As outlined in the following table, the price per GB typically correlates to the number of program/erase cycles a flash device is rated for.

2. FLASH COMES IN DIFFERENT FLAVORS

	SLC	eMLC	cMLC	TLC
Flash Cell P/E Cycles	100,000	10,000	3,000	300
Performance	Highest	High	High	Medium
Price per GB	Highest	High	Moderate	Low
Bit Error Rate	Low	Low	Low	High

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

SLC (single-level cell): Once considered the only form of flash appropriate for enterprise storage, this format of flash stores one bit in each cell. As a result, it has high reliability and cost. SLC is no longer being fabricated, in part due to the success of all-flash system designs that have significantly raised the reliability of eMLC and cMLC flash.

EMLC (enterprise multi-level cell): MLC increases storage density and reduces cost per GB by storing two bits per cell. eMLC is a form of MLC that increases the P/E cycles by reserving a significant portion of their total flash address space for over provisioning (OP), a percentage of the total physical memory reserved (commonly 28% to 50%) to replace cells as they wear out. eMLC flash often includes advanced controller firmware to provide device level optimizations like garbage collection (GC).

MLC (consumer multi-level cell): MLC increases the storage density and reduces cost per GB by storing two bits per cell. cMLC is the form of MLC used in the vast majority of consumer devices and of late, enterprise storage arrays. cMLC is the highest volume form of flash being fabricated today. cMLC includes a low amount of over provisioned flash (7% of address space) and a limited set of capabilities in controller firmware.

TLC (triple-level cell): The most recent innovation in the flash market, TLC allows for 3 bits to be stored in each cell, providing a means for flash to better address capacity centric storage needs. The architecture is more susceptible to bit errors than MLC due to the narrow band in which the bit's value can be read. In addition, TLC supports hundreds of write cycles, meaning it may wear out quicker than MLC.

3. FLASH FAILS DIFFERENTLY THAN HARD DISKS

Both flash and hard disk drive increase in failure rates over time; however, why and how each fail is significantly different. 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 hard disk.

- Hard disks experience mechanical failures due to vibration, voltage regulation and excessive heat. Media bit errors are often with high adjacency – think of a scratch on an LP (you do recall the 331/3 rpm microgroove vinyl record right?)
- Flash experiences bit errors based on wear (PE cycles) – the media simply gets worn over time. Access degrades over time, first to a read-only state followed later by complete device failure. Further, the unrecoverable bit error rate (UBER) 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 hard 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.

4 CLASSES OF ALL-FLASH PRODUCTS

PCIE CARDS IN SERVERS	FLASH APPLIANCES	RETROFITTED ARRAYS	ALL-FLASH ARRAYS
<p>PURPOSE: Acceleration of a single application</p> <p>KEY ATTRIBUTES:</p> <ul style="list-style-type: none">• Flexible use as either an IO acceleration cache or as persistent storage• No high availability: must mirror datasets, doubling cost of flash• Relative low capacity 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• Capacity: 1s of TBs	<p>PURPOSE: Acceleration of a single, multi-host application</p> <p>KEY ATTRIBUTES:</p> <ul style="list-style-type: none">• Combines very fast performance with high capacity• JBOF: just a bunch of flash• Typically feature limited – no enterprise storage features (NDU, HA, snapshots, replication, etc.)• Dense / compact per RU• Capacity: 10s of TBs	<p>PURPOSE: Acceleration of limited set of applications</p> <p>KEY ATTRIBUTES:</p> <ul style="list-style-type: none">• Existing storage array retrofitted with all-flash configuration• Increases performance while keeping similar management model for existing customers• Complex setup and deep expertise needed for ongoing operations• Tradeoffs are inherent across performance, resiliency, efficiency, and simplicity• Disruptive operations such as data migrations and downtime• Weak or no data reduction leads to high \$/GB• Customers typically need to re-purchase hardware and software every 3-5 years• Capacity: 100s of TBs	<p>PURPOSE: Next-gen tier-1 storage infrastructure</p> <p>KEY ATTRIBUTES:</p> <ul style="list-style-type: none">• High capacity, shared storage with enterprise-class features• Not the absolute fastest, but delivers sub-millisecond IO access• High availability, non-disruptive upgrades, snapshots, and other enterprise features included in software stack• Handles mixed workloads efficiently• Highly efficient, reducing the physical and energy footprint through built-in data reduction including deduplication and compression.• Next-generation simplicity and support model• Capacity: 100s of TBs

WHAT ALL-FLASH STORAGE CAN DO FOR YOU

Before we proceed, let's pause to consider the applicability of all-flash storage and set expectations around the benefits one should expect when modernizing your storage infrastructure. Flash can pave the way for you to gain numerous benefits for your applications – not just those that demand the highest levels of performance.

1. GREATER APPLICATION PERFORMANCE

Flash easily delivers 10X the performance of hard disk drives allowing you to accelerate your applications whether that is processing more data transactions per second or dramatically reducing the time required to run batch processes. The performance promise is where many begin their journey to adopt flash.

2. RESPONSIVENESS FOR A DIGITAL ECONOMY

The real advantage of flash storage is consistent sub-millisecond latency, making every application more responsive. Reduced latency has been cited by Amazon, Microsoft, Google and Walmart as the key to prolonging customer engagement, increasing sales, and raising user satisfaction in today's digital economy.

3. OPERATIONAL SIMPLICITY

Flash can greatly simplify application administration and storage operations by greatly reducing, and in many cases eliminating, configuration tuning and tweaking. Adopters of flash reclaim hours of staff time that can be reinvested and applied to strategic initiatives that advance the business.

4. ACCELERATE NEXT-GEN INITIATIVES

CIOs seek to gain a competitive advantage with initiatives that often include agility, cloud (and cloud operating models), security, mobile, big data analytics, and the Internet of Things (IoT). Flash provides the foundation that ensures next-gen initiatives will run as well in production as they do during the pilot and testing phases. The economics and agility of flash also accelerate the rise of internal clouds as a more secure, more performant, more resilient, and more cost effective alternative to the public cloud – giving enterprises greater flexibility and control around their business critical applications.

5. DATA CENTER EFFICIENCY

Flash consumes a fraction of the power and rack space compared to disk storage. By adopting flash, data centers achieve resource efficiencies that allow them to host more IT services and store more data well into the future.

EXPECTATIONS FOR FLASH

With flash prices now below performance disk and still continuing to aggressively decline, you should consider all-flash storage options as the means to modernize the tier-1 storage infrastructure today and for capacity centric data sets or tier-2 storage in the not too distant future.

In order for these benefits to be realized, you must consider how flash is implemented, and the benefits and the trade-offs associated with each of the four forms of flash. The subsequent sections will dive into common storage platform design goals, specifically availability, resiliency, affordability and operational simplicity.

AVAILABILITY: IS THE ARCHITECTURE DESIGNED FOR NO DOWNTIME?

Often storage conversations fall on architectural lines without defining business goals or how the benefits of the architecture are delivered. Consider storage arrays, where the conversations tend to fall into scale-up versus scale-out architectures. The storage architecture must be designed for flash in order to unlock and free one to capitalize on the full capabilities of flash memory.

While a traditional storage array can be accelerated by the addition of flash, tradeoffs and consequences may arise during times of controller failure and maintenance. Understanding how these operational scenarios are remedied is something to know before you buy a product.

Is the Controller a Stateless Architecture?

Modern applications, virtual and cloud infrastructures and even blade servers leverage a stateless architecture to separate application service delivery from the underlying infrastructure. Wouldn't this be ideal with storage?

The majority of storage platforms implement an architecture that binds the processing of IO with the persistent storage medium. This model leads to inconsistent performance during hardware failures and maintenance operations and quite often requires data migrations with hardware upgrades.

A stateless storage controller architecture separates the IO processing plane from the persistent data storage plane, providing high availability with non-disruptive operations. This architecture enables hardware platform updates, software upgrades and storage expansion without reconfiguring applications, hosts, or IO networks – all without disruption or loss of performance.

Is Performance Impacted During a Failure or Maintenance?

We all know mission critical applications cannot tolerate outages, however the adoption of flash raises the stakes. The downside to flash storage and sub-millisecond latency is the impact on an application due to performance loss with a failure or maintenance event.

A flash architecture must reserve CPU and memory resources to provide the same consistent performance during periods of failures and maintenance as during normal operations. A 5X or 10X increase in latency could be catastrophic to an application or business. We recommend you understand how performance is delivered when the inevitable failure occurs.

Is Performance Impacted with Array Utilization Reaching 100%?

It's important to understand whether the flash array can deliver full performance as the array fills up to 100% full – and even beyond. Some arrays do not offer sustainable performance, as a performance drop-off will be visible after 70-80% utilization. This reduces the effective usable capacity of the array, thus reducing ROI.

How Invisible is a Non-Disruptive Upgrade?

Most all-flash storage solutions support some form of a non-disruptive upgrade today. Some require proactive planning or days of background data migration, while others can scale only as long as the hardware is identical – locking customers in to quickly outdated hardware models.

Determining whether it is possible to replace a controller or a component without having to take the flash system offline is just the start. Does the upgrade require applications to be adjusted, host volume managers and multi-pathing software reconfigured? What about the network, does the expansion also require additional IO ports to accommodate the new hardware?

How long does all of this take? Wouldn't it be ideal if the flash architecture simply allowed for hardware and software upgrades to be completed in a manner that was non-disruptive and completely invisible to the surrounding infrastructure?

RESILIENCY: HOW IS THE DATA PROTECTED?

The attributes of flash memory, particularly its performance and infinite PE cycles, requires you to reconsider data protection mechanisms. With disk storage, you had to complete complex calculations with considerations for host IO requirements, write penalties, media errors, and rebuild times that correlated with capacity in order to hope to match performance and risk mitigation goals. While the performance of flash will reduce rebuild times, there's much more to consider before protecting it like it's disk.



How Many Drives Can Fail at Once?

You need to understand the number of concurrent failures a data protection scheme can support and the cost of data protection overhead, whether its RAID, mirroring or erasure encoding. While flash will dramatically marginalize the performance impact of failures and reduce the associated rebuild times – speed is not a substitute for reduced data protection. As capacities for both disk and flash storage devices are measured in terabytes (TBs), dual parity RAID schemes have become the de-facto means for enterprise class data protection against unrecoverable bit errors and drive failures.

How Much Overhead is Required for Data Protection?

Data protection is critical, however both parity based RAID schemes and mirrored data replicas consume storage capacity. Understanding the overhead associated with data protection can help you better understand the true effect on the cost of storage.

If you're considering a server-based (with PCIe cards) or any shared-nothing architecture, you'll be required to reserve enough capacity to support the number of device failures in the protection scheme. This goes beyond a simple mirror and could include an overhead as high as 67% – in a 3 node configuration that supports two concurrent failures. This design results in a maximum hardware efficiency of 33%.

Dual-parity RAID schemes designed for flash can provide high data protection with a minimal capacity overhead – often around 20%. These flash-optimized designs eliminate the multiple IO overheads associated with storing and updating parity in RAID designed for disk arrays. You should steer clear of RAID protection schemes designed for hard disk drives as they will accelerate the PE cycles on an SSD.

Support for Snapshot Backups and Data Replication?

Up until now, this section has focused on data protection in the event of a system or component failure but as you may know, most data loss is the direct result of human error.

A key enterprise feature of shared Tier 1 storage is the ability to take local snapshot backups and to replicate data to remote sites for the purpose of disaster recovery and on-site backup. With flash memory these processes are lightning fast. However, due to PE cycle concerns, you must assure that the snapshot mechanism is based on metadata and does not copy the data out as data is updated or deleted.

For most solutions, replication is a native feature -- however it is not a universal capability especially with PCIe flash cards and flash appliances. Third party replication options provide data replication with support for heterogeneous environments. These solutions add to the total cost of storage as they include additional costs for software licensing, network ports, and replication hardware.

AFFORDABILITY: DATA REDUCTION, SOFTWARE LICENSING AND MAINTENANCE COSTS

Economics play a key role in whether you can adopt flash as their next-gen storage infrastructure. For some, cost comparisons are as simple as comparing acquisition prices – valuing the price per raw GB of storage capacity without consideration for effective capacity, software licensing, the inevitable rise in maintenance and support costs, and the cost of upgrades and migrations to subsequent product generations. You need to know the total cost of flash over a reasonable timeframe (typically 5 or 6 years) or you may be in for a big surprise.

Data Reduction: Amplifying Effective Storage Capacity

You will find a number of different methods of data reduction in the market around flash storage – these should not be treated as quantitative lists of items to check off. The means by which data reduction is implemented will have a material effect on the performance, efficiency and cost-per-GB of effective storage.

Data reduction may sound like a commodity but implementations are as diverse as the number of flash options within the market place. It's important to understand the likely returns on these technologies as they will make or break your budget. The basic rule of thumb is that the more comprehensive the data reduction technology, the greater the savings capabilities - and the greater the applicability of the solution to your environment.

Defining Common Data Reduction Technologies

Below is a short primer on the most common data reduction technologies in the market today:

- **DATA DEDUPLICATION:** Ensures only unique data is stored on a storage array/platform. This minimizes capacity consumed and eliminates the need to write redundant data.
- **DATA COMPRESSION:** Reduces the capacity required to store unique data on a storage array/platform. This minimizes capacity and number of writes required to store unique data.

Be aware that while some vendors include Thin Provisioning in their data reduction savings claims, Thin Provisioning is not a data reduction technology; rather, it is a dynamic means to allocate storage capacity on demand.

Global vs. Optional Data Reduction and Performance

Storage architectures retrofitted with flash often make data reduction optional due to the negative impact on performance.

Optional data reduction technologies are often marketed as providing customers the choice between maximum performance or capacity optimization. Such claims fail to include the performance trade-off that makes these features optional. Buyer beware when selecting optional storage features.

Inline or Post-Process Data Deduplication?

Some vendors will reduce data as it is coming into the array, others will first land data on flash before processing for data reduction. For traditional disk arrays, post-process deduplication was the norm because data reduction took a long time to execute and could be scheduled to minimize the performance impact. Scheduling data deduplication may cause a loss in capacity savings when not timed correctly due to data being locked in snapshots and replication processes.

With flash, inline deduplication is not only possible due to the performance of flash and new architectures designed for it, but inline deduplication is critical in raising the reliability of flash by avoiding writes and eliminating PE cycles for data that already exists. Inline eliminates the challenges historically associated with disk storage.

Fixed vs. Variable Length Deduplication and Granularity

Data deduplication implementations vary in the block size granularity at which they analyze data with most falling into one of two categories based on the storage platform's back end architecture – they are either fixed or variable in length. The smaller the granularity of the block size used to analyze data, the more redundancy and capacity savings will be returned.

VARIABLE LENGTH DEDUPLICATION: provides the greatest level of granularity, allowing redundancy that would normally not be identified with fixed block implementation to be identified. This model requires more CPU cycles but returns the greatest possible data reduction.

FIXED LENGTH DEDUPLICATION: limits the granularity of the block size to match that of the backend. This form of implementation optimizes performance by limiting the number of CPU cycles dedicated to identifying redundancy at the expense of lower reduction levels.

How Many Compression Algorithms?

Compression provides a varying range of data reduction based on the data set. Some data respond better to one form than another and some respond well to a combination. Flash storage that offers multiple data compression algorithms can be relied upon to provide a quantifiable advantage in data reduction over those with only a single form.

What About Host-Based or Storage Gateways for Data Reduction?

Data reduction can be delivered in host applications and via storage gateways which work well for host flash, flash appliances, and flash retrofitted disk arrays. By requiring additional software licensing, network ports and hardware appliances these solutions increase the total cost of storage – which is counter to the primary goal of native data reduction.

Native data reduction capabilities in the storage infrastructure not only eliminate the additional cost and points of management but also provide significantly better reduction results.

Is the Array Software Included or is it an Additional Cost?

Traditional storage vendors often license their data service software (protocols, snapshots, replication, cloning, etc.) based on capacity, which often requires customers to purchase new, more expensive versions as their data footprint grows or they upgrade their storage hardware.

This model seems archaic and out of date with the rampant rates at which data is growing globally. Flash buyers need to fully understand the cost of storage software associated with a given architecture.

Do Maintenance and Support Costs Increase Over Time?

The traditional model for storage maintenance and support needs to be modernized. Historically speaking, storage vendors dramatically increase these costs once the initial support contract expires. The increase in costs was used as an incentive – or stick based on one's vantage point – to entice customers to upgrade to newer hardware at the same or lower cost. This process mostly benefits vendor sales while forcing customers to take on lengthy and complex data migrations.

Does your flash vendor offer a flat maintenance renewal program? Such new models have emerged and offer a modern means to plan for future storage costs.

What's the Infrastructure Tax When Scaling?

Some storage arrays scale performance and capacity independently while others require one to scale in fixed, hardware-defined units that require additional power supplies and multiple network IO ports. Lastly, server based scaling requires redundancy overhead to protect data and provide HA.

What's the Density of the Infrastructure?

Beyond just the capacity offered by a storage array, consider the density of the storage. The denser the storage, the smaller the physical footprint and more precious space freed up in the data center.

OPERATIONAL SIMPLICITY: HOW TO SCALE THE BUSINESS

Providing storage services today is a complicated and time-consuming task at best. Even as gains are made in terms of application integrations and automation, the storage team is often still lost in a sea of complexity. Flash will make the environment simpler – but how much simpler requires a bit of knowledge.

How Much Tuning Does the Flash Require?

Flash should be simple – so what defines simple? Is a next-gen flash solution more simple than retrofitting a traditional disk storage array? It's hard to debate your comfort level for what you know; however, familiarity does not equate to an increase in productivity.

Significant advancements can be found in arrays designed specifically for flash that simplify operations while increasing reliability and performance. Most of these platforms have simplified storage down to provisioning based on capacity and establishing data protection policies (for snapshot backups and remote replication).

Ask about how flash can eliminate antiquated operational processes like defining storage pools and RAID configurations, optional settings around cache, tiering, data reduction, etc.

How Much Application Tuning Does the Flash Require?

Even if you have flash you may still need to tune your application to achieve optimal results. For flash platforms with fixed block sizes, applications may need to be exported and imported to align to the IO transfer size – this process is a common recommendation with OLTP databases, VDI and virtual infrastructures.

Do you have to load balance your application across volumes and/or controller nodes? Some flash solutions limit performance on a per-volume or per-node basis and require considerations for application data layout – both today and in the future as an application grows.

If you are considering deploying PCIe flash in a host, you will likely have to consider if the flash is being used as an acceleration cache or as persistent data storage.

What Application Integration Exists Today?

Storage integration capabilities exist to deliver resources with agility in an on-demand manner. Such integrations can provide organizations with a level of on-demand service that once was only available in public cloud offerings. Native support for application integrations should be preferred over 3rd party integrations based on cost and complexity.

Does the Flash Provide a Programmable Set of APIs?

DevOps and the desire to further the integration of applications and workflow can be realized through programmable API interfaces. REST based frameworks tend to make a solid foundation and allow for the extension or layering of additional scripting languages like Python, Perl and the C library. For Microsoft environments, you will want to ensure the inclusion of an SMI-S API library.

What Management Tools are Available for Flash?

Ongoing management simplicity plays a key role in minimizing operational costs and reducing the risk of unplanned downtime. Although flash provides a solid foundation, some products offer an outdated management model on top of that flash, bringing scalability limitations or complexity that hampers the agility – hampering the strategic goal of bringing a cloud IT model to internal infrastructure. Next generation management models are built using a SaaS approach, eliminating the need to deploy and manage still more internal infrastructure. SaaS approaches also ensure that you get access to the latest management and analytics features automatically. Therefore, as you consider all-flash storage, also evaluate the management model and determine how easy it is to consume, and to keep current as vendors introduce new innovations.

What Support Model Does the Vendor Offer?

Undoubtedly, outstanding support is needed to address issues quickly and maintain demanding service level agreements. However, the support provided by all-flash vendors varies considerably. In its more primitive form, support is reactive – with customers filing tickets and waiting for support to resolve incidents. However, next generation support combines the power of analytics to collect infrastructure data and predict scenarios that can lead to issues even before they arise. Hence, proactive and predictive support can make a strong difference and enable you to maintain, and exceed, service level agreements.

HOW MATURE IS THE FLASH 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 – even though they are advertised. Understanding whether you are the first customer ever to lay hands on the product and how ready it is for your specific environment will help you make a good choice.

Multiple releases of a product – both the hardware and software – can represent a level of maturity and success in the market. and should be considered over 1.0 product releases and promises of future enhancements.

Benchmarks and Certifications

Unfortunately, many of the benchmarks in the market are designed to measure the performance of mechanical performance disk. Therefore, 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 with real-world application workloads 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.



Are Analysts, Customers, and Competitors Sharing Their Opinions?

Where can you turn when you need a third party validation of a vendor's product? The answer is to look for external validation – it's often worth much more than vendor claims around their own product.

- **INDUSTRY ANALYSTS:** While often not the most aggressive around the adoption rate of emerging technologies, it appears that industry analysts are now acknowledging flash storage as the new form of Tier-1 storage. There are a number of reports that rate the vendors, their products and technologies and should be highly considered in your selection criteria.
- **CUSTOMER EXPERIENCE:** Surveys such as the Net Promoter Score (NPS) collect broad-based customer experience surveys and translate them into a vendor and product rating. While customer experience polls may not provide the final decision criteria, they should be used to rule out those who under deliver from a customer experience standpoint. Also, NPS ratings should be validated and certified by third parties such as Satmetrix, to bring transparency to the survey methodology and results.
- **COMPETITIVE TARGETING:** This may be the least likely example of a product endorsement but if you want to know which products are truly delivering benefits and new capabilities in the market, consider looking for trends amongst market competitors. If two or more incumbents have proactive negative statements regarding an upstart – chances are the upstart has something the incumbents are lacking.

CONCLUSION

Modernizing your data center infrastructure does not occur often. Such strategic investments must be designed to support a broad set of use cases as they are amortized over a long period of time.

Flash storage can pave the way for your organization to provide greater scale to applications, IT operations and overall data center capacity. Flash adoption carries the promise to transform organizations and businesses, allowing greater insight into information, streamlining operations and accelerating the delivery of new and innovative IT services.

In determining the best flash storage system and flash storage systems vendor, you must ask the right questions not only of the vendor, but about the application needs and the areas requiring additional performance. Using the criteria laid out in this guide, it will help you determine where and how flash storage can help improve the performance and agility for your enterprise IT infrastructure.

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



sales@purestorage.com | 800-379-PURE | @purestorage