





Invest in the right
**flash storage
solution**

A GUIDE FOR THE SAVVY TECH BUYER

Introduction

You're looking at flash storage because you see it's taking the storage world by storm. You're interested in accelerating business-critical applications, consolidating a virtual server or desktop deployment, trying to get ahead of your company's data onslaught, or some combination of the above.

By 2019, 20% of traditional high-end storage arrays will be replaced by dedicated solid-state arrays (SSAs).¹



This easy-to-read guide was developed to help arm you with key considerations and questions to ask before investing in a flash storage array for your business today, and for the future.

What are the key considerations for flash buyers?

1. Can I afford flash?
2. Is it competitive?
3. What do I need to know beyond the flash hype?
4. What questions do I ask vendors?
5. What are the risks?
6. What about software vs. hardware?
7. What are some of the "gotchas"?
8. What do I need for the future?

¹ Gartner RAS Core Research Note G00260420 Magic Quadrant for Solid-State Arrays, Valdis Filks et al, 28 August 2014.

Can I afford flash?



High performance



More data density



Cost savings

You've heard it time and time again, "We want higher performance and we've got to store more data with less resources, at the lowest possible cost." Flash is a no-brainer answer to the performance requirements, however, until recently flash vendors could not compete with spinning disk prices.

Affordability—what to look for

Data compaction technologies

Look for technologies such as deduplication that can reduce overall capacity requirements and reduce TCO. Watch out for compaction ratio promises. Some will be better and some will be worse, just know that your dollar per GB can come in higher than the vendor promises.

Drive capacity and system density

Consider utilizing high-density flash drives (e.g., 7.68TB or 15.32TB SSDs) which offer space benefits. Data compaction technology in combination with large capacity flash can reduce footprint by 80 percent or more, while lowering power consumption commensurately.

Moving live data

Moving aging data from flash to lower cost media can be smart and cost effective. Whether you choose a converged flash array that offers spinning disk alongside flash media or software to move live data from an all-flash array to a disk or hybrid array, both can improve your TCO.



\$1.20/GB

You can now get all-flash storage for as low as **\$1.20 per GB useable**. Less than half a rack of flash array can replace four racks of traditional high-end storage.

Is performance the same for every flash array?

The performance bottleneck has shifted from the drives within the array, to the ability of the controllers to perform (cores, memory) and scale (multi-controller). When it comes to performance, flash arrays vary widely, not just in terms of performance, but in terms of consistency and predictability of performance.

For example, workloads require not only high performance (in terms of IOPS) but also depend heavily on predictable, low latency response times. Flash media injects a whole new set of requirements that may not be addressed on all-flash arrays. Simply adding flash drives to a traditional disk system will not provide consistently high levels of performance and predictable low latency. That's why a flash-optimized architecture is so important.

Whether you choose an all-flash, hybrid, or converged flash option, it's important to look for an architecture that is designed to meet service levels for a diverse range of applications.



Workloads require not only high performance, but fast response times.



Did you know there are hybrid, converged, and all-flash arrays?



Hybrid arrays

Lower in cost than an all-flash array, a hybrid storage array with auto-tier capabilities can be a solid option for accelerating applications with minimal incremental investment. By employing a flash tier in addition to spinning media, hybrid arrays can significantly accelerate performance.



All-flash arrays

All-flash arrays are composed exclusively of solid state drives. When performance matters most, nothing can beat an all-flash array built on a flash-optimized architecture. Look for an all-flash array that can deliver consistent performance without compromising resiliency, scalability, or data mobility.



Converged flash arrays

A new concept in the flash world, the converged flash array is a category of all-flash array that can accommodate disk media as well, offering the performance and latency of an all-flash array, the affordability of a hybrid array, the agility of a unified array (with support for block workloads, file shares, and object access), and the scale and resiliency of a high-end array.

Performance—what to look for

Scalable, consistent, and predictable performance

Look for flash solutions that employ architectural elements that enable CPU off-load and capabilities such as quality of service (QoS). Architectures that use dedicated hardware to perform specific functions, such as RAID parity calculations and deduplication, can free up (offload) array CPU resources. In addition, flash media is so fast that it can easily saturate storage controllers, so aim for solutions that scale beyond 2-node architectures to enable greater performance and capacity scalability within the same system.

Deduplication only where you need it

To help optimize performance requirements and capacity efficiency requirements at the same time and without service interruption, you may want to consider enabling and disabling deduplication on a per-volume basis depending on the data set.

IOPS, latency, and bandwidth

IOPS, latency (time taken to service an I/O), and bandwidth (overall data throughput available) are basics of any flash array (or even basics of any storage system). Don't just look at single metrics. Make sure you look at the IOPS at a given latency level. For example, a solution claiming to deliver hundreds of thousands of IOPS is of little use if the associated latency is not sub-millisecond consistently.

Performance during failures

Component failures happen, so consider massively parallelized architectures. They tend to perform very well even under failure conditions and are designed to easily pick up any extra load from a component failure.

Flash optimized architecture

Be aware that flash media is fast and can easily saturate every other component in the system (controllers, ports, I/O inter-connects, etc.). A flash optimized architecture ensures that every component in the system is utilized. For example, imagine each virtual machine has every SSD, controller, and port at its disposal for any workload spike—this is the benefit of such an architecture.

Why keep talking about resiliency and scale?

You've heard it over and over, "The always-on, from-anywhere world creates exponentially more data every year." You know what that can look like. A robust platform with plenty of "headroom" today may not be adequate in a couple years as data (and cost of downtime) grows. So beware!

It's important to know that when it comes to flash arrays, enterprise-class resiliency and ability to scale are not always a given. There's no point in having the power of flash if you are out of available capacity or it's not available 24x7, 365 days a year.

Plus, enterprise-class features like synchronous/asynchronous replication, live data migration, and transparent data center failover solutions might seem like table stakes these days in tier-1 storage, but not so in the new world of all-flash arrays. Many "new" all-flash arrays lack these enterprise-class data services. Some deliver them using external appliances, and that only adds cost, complexity, and downtime.

If you are going to put your most important data on an all-flash array, you must have a way to access, manage, and scale it at all times.




When it comes to flash arrays, enterprise-class resiliency and ability to scale are not always a given.



Consider platforms that scale into the 100s of terabytes.

100 terabytes



Resiliency and scale—what to look for

Proven enterprise-class

If you plan to run production workloads on your array and need to support mission-critical applications, your array must be capable of delivering 99.9999 percent availability.

Bulletproof your data

Look for an array that offers hardware and software redundancy, advanced high availability features like non-disruptive hardware and software upgrades (be sure to ask about multi-pathing dependencies—best to avoid them altogether), data integrity with T10 DIF, and transparent failover solutions.

Data center test feature set


For example, remote synchronous/asynchronous replication, application integrations (with Oracle, SQL, VMware®, Hyper-V, etc.), and non-disruptive data mobility between systems are critical but not yet standard offerings in the all-flash world. You might be surprised by the number of solutions that lack these key capabilities.

Modern scale

Given the average midrange system is about 150 TBs and high end 400 TBs+, a flash system supporting 30 or 40 TB is hardly a capable replacement. Consider platforms that scale into the 100s of terabytes or more.

What about backup/recovery?

The availability of your data is critical to your business operations. On those rare occasions when system failures occur, you need to be ready and able to provide uninterrupted service. That means data must be efficiently backed up and immediately recoverable, without adding unnecessary overhead or cost to day-to-day operations.



Deduplication can reduce backup storage by a factor of 20.



Application-integrated data protection



Simple and efficient backups



Reduced risk exposure

Data backup & recovery—what to look for

Application-integrated data protection

Look for storage arrays that integrate with the tools you already use, enabling hypervisor and application owners to control backup and recovery processes directly from their preferred system management consoles. For example, if you're running VMware, look for a solution that will allow you to manage snapshots and initiate recovery from within the VMware vCenter™ management console. The best solutions will also offer the ability to create application consistent backups of leading business applications. Even better if there is an SDK to enable partners and developers to integrate with their own business applications.

Simple and efficient backups

An important component of backup processing is deduplication, which can reduce your backup storage requirement by a factor of 20. Scalable, non-duplicated backups can also enable a greater number of recovery points, giving you more granularity and control. Find out if deduplication is native to the storage solution or requires server resources. Be sure to also ask about the performance impact of snapshot generation on application servers—it should be minimal or non-existent. Look for tools that make it simple to automate and manage these processes.

Reduced risk exposure

You can reduce your exposure to risk if your flash storage solution can deliver on the same stringent recovery SLAs as your current disk-based environment. Not all flash systems can do this. Look for a solution that creates fully independent backup volumes that can be restored at the volume level in the event of disaster. This level of data protection cannot be achieved with snapshots alone.

Just as there are differences in cost and performance between flash arrays, there can also be significant differences in the cost and complexity of providing and managing data protection for flash storage arrays.

Your business operation is dependent on the availability of critical data. Backup processes that don't fully protect your data leave you vulnerable. Backup should be non-intrusive. It should never create performance impacts on production servers, nor should it require system downtime. Recovery processes should be simple and efficient. Data protection shouldn't require expensive and complicated backup appliances or software.

If you are going to put your most important data on an all-flash array, it must be fully protected and easily recoverable. Anything less puts your most critical business asset at risk.



Your most important flash data must be fully protected and easily recoverable.

What you now know

You now know that flash is affordable and finally competitive with spinning disk at \$1.20 per GB or lower pricing. You know that not all flash arrays are equal. You know that flash can save up to 80 percent space or more compared to HDD-based solutions. You know that a true enterprise-class solution can give you 99.9999 percent availability even under failure conditions. You know to look out for an enterprise-class solution for scale and resiliency, and that data mobility might be important.

Important questions to ask your vendor(s):

Affordability

1. Do you have data compaction technologies? How do they work?
2. Which generation of flash drive are you using? What's the capacity?
3. How does the cost compare to your current \$/GB—raw and usable?



Performance

1. Does the array offer consistent, predictable performance with sub-millisecond latency?
2. How does the array perform under failure (drive, controller, etc.)?
3. What kind of QoS controls do you provide to manage my most important workloads? Can they help me deal with “noisy neighbors”?
4. Does it offer QoS controls to support multi-tenancy and consolidation?



Resiliency & scale

1. Do you have high availability features? What are they?
2. Will the solution scale sufficiently to be deployed where I used to deploy midrange or high-end arrays?
3. How resilient is the architecture? Please describe the capabilities.
4. What disaster recovery options does it offer?




Data protection

1. What backup/recovery capabilities are available?
2. How will backup/recovery integrate with my current environment—can I manage from existing hypervisor or applications?
3. Will I be able to match or exceed current SLAs?
4. What is your approach to deduplication in backup and recovery processing?
5. Are there performance impacts or downtime associated with backup/recovery processing?
6. How do you ensure recoverability of application data?



Other considerations

1. How well will it integrate into my existing infrastructure?
2. How easy is it to manage?
3. How easy is it to migrate data on and off my new array?
4. How does the array use technologies to extend media lifespan?
5. Does it support file shares and object access?
6. Will the vendor be there for the long haul?





Want to learn more about flash-optimized data protection?

Learn the four key considerations



Sign up for updates