

Technical White Paper

All-flash storage is breaking these 5 storage rules



Contents

- Evolution of Storage..... 1
- Flash is Changing Five Core Storage Rules 1
- Rule 1: Controller determines flash array performance 2
- Rule 2: Data reduction is table stakes..... 2
- Rule 3: Integration with cloud management orchestration platforms..... 3
- Rule 4: Commodity hardware rules productivity curve..... 3
- Rule 5: Software-defined innovation requires flash performance..... 3
- Tintri: Advancing the 5 core storage rules with VAS 4
- Conclusion 5


Evolution of Storage

It should come as no surprise to anyone that storage platforms have seen accelerated innovation over the past 15 years. Deduplication, in-line compression, dramatic cost/GB reductions across technologies, exponential performance growth, specialized storage and cloud have all turned traditional storage status quo on its head. One of the latest and most significant changes has been the broader adoption of flash.

Flash presents a storage platform that provides a quantum improvement on the value storage has on an organization's agility. Originally limited to accelerating HDD enterprise arrays, improvements in its cost and performance have transformed the storage industry and led to the development of economically viable all-flash arrays (AFA). This paper will investigate five critical considerations in choosing an AFA and how VM-aware flash from Tintri changes the game in one big use case: optimizing your virtualization platform.

Flash is Changing Five Core Storage Rules


Like any new game-changing technology, one of the biggest inhibitors to flash's adoption has been its cost. In the early days of flash, accessing its 10x advantages in performance, power and density meant an associated 10x increment in cost/GB. This effectively relegated flash to very high-end applications and only where absolutely required, such as supercomputing or complex algorithm or video processing. In a very short span of time, however, flash has seen a dramatic 80% cost reduction and its advantages have accelerated in terms of performance, power and density. This perfect storm of innovation has pushed flash deeper into traditional enterprises and increased its adoption rate.



Today, AFAs have become a more feasible option. As the cost of flash falls and OSs are optimized further to leverage the speed of underlying hardware, AFAs present a high-performance storage alternative with compelling economics.

For virtualization and the cloud in particular, high-performance applications have been challenged even further to deal with HDD limitations, especially around performance and latency. AFAs provide the performance and latency capabilities needed by Tier-1 and Tier-0 applications but face other ancillary challenges.

One of the key benefits of virtualization is the flexibility it provides for dynamic allocation and placement of VMs through vMotion and DRS. But this dynamic reconfiguration of sharing by VMs of LUNs/volumes generates the I/O blender effect, performance problems and contention for storage resources. flash addresses these issues to some extent but the underlying challenges of mappings, LUNs/volumes and similar legacy constructs remain.



As the evolution toward flash as an independent storage tier continues, Tintri summarized countless hours interacting with the market into five key rules affecting this new paradigm.

Rule 1: Controller determines flash array performance

To scale an HDD storage platform, the traditional approach is adding more, more, more. Need more capacity? Add more shelves. More IOPS/ performance? More shelves (expanding striping). But more shelves equals more complexity, more cost, more power/space and more chance of failures.

Because flash boasts high performance and low latency, it eliminates once-necessary data placement tricks from the controller. Meanwhile, much of the bandwidth of an HDD controller is consumed managing reads and writes to overcome or mask latency inherent in the technology. This is also a reason for sub-optimal flash performance when it is simply bolted on to a traditional HDD architecture. Flash-optimized architectures avoid this pitfall by jettisoning legacy HDD structures.

Yet flash can provide significantly greater levels of performance than is realized using just a standard flash controller, resulting in the controller becoming the bottleneck unless it is highly optimized for the array. This means that adding shelves to AFAs does deliver more capacity but may not deliver a commensurate increase in performance/IOPs.

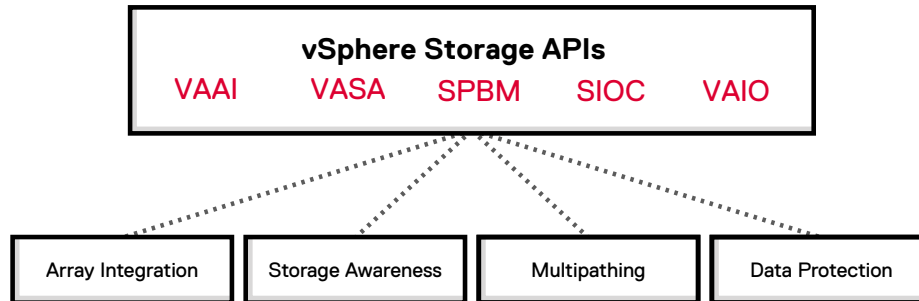
Rule 2: Data reduction is table stakes

When flash first started its adoption cycle, it was not common to see data reduction techniques applied to mission-critical storage. This is because in-line data reduction services like deduplication, compression and cloning/thin provisioning had a significant effect on latency for HDDs, so these services were relegated to out-of-band storage functions like backup and archiving.

Because of the near-instantaneous responsiveness of flash, however, data reduction comes free and is expected now. This has resulted in these inline data reduction techniques being essentially mandatory for AFAs today.

Rule 3: Integration with cloud management orchestration platforms

The best data reduction for a virtual system is achieved through solid integration with the hypervisor, such as with vSphere and vRealize. This can be as basic as supporting VAAI, VMware's standard storage API for array offloading, to avoid significant levels of data traffic. Other useful VMware integrations include VASA, storage policy-based management (SPBM), storage I/O control (SIOC) and VAIO. As a rule of thumb, the more of these integrations a vendor supports, the greater the cost savings for users.



As integration with multiple cloud orchestration platforms become critical, look for a vendor that has done the integration and is on the HCL to set the stage for achieving higher levels of performance and data reduction from flash.

Rule 4: Commodity hardware rules productivity curve

Initially, flash was deployed either as stand-alone SSDs on a local host or as proprietary hardware such as a PCI-E card. The former gave good local performance but didn't provide enterprise class storage. The latter provided high-end performance but at the cost of reliability.

Today flash architectures have coalesced around two different paths. One is built on custom hardware and interconnects based on flash to squeeze every bit of performance possible out of the technology but typically at a significant price premium. The second is built on a platform that can ride the cost curve of commodity flash and off-the-shelf hardware supported by robust architectures and data services, much like how custom compute platforms have standardized on commodity Intel- or ARM-based processors.

As the cost/GB of flash falls under the cost curve for HDD, this second option makes dramatically better economic sense for all but the most specialized and/or esoteric workloads (HPC, big data modeling, high-speed, real-time custom kernel workloads). Commodity-based flash platforms are a much more cost-effective choice for workloads like server virtualization, VDI, cloud applications and object and file storage.

Rule 5: Software-defined innovation requires flash performance

The performance and latency benefits of flash can dramatically improve innovation at the software layer, something that can be missed if the only goal of flash is as a high-performance data repository. Simply bolting on flash as a cache to an existing HDD-based platform can provide some level of performance acceleration but at significant economic and opportunity cost (missing out on all the additional capabilities/benefits outlined in this paper). This significantly bypasses the improvements possible for virtualized workloads with flash.

Historically, focus and differentiation was around hardware innovation built on commodity drives. Custom ASICs and propriety controllers were the foundation for low-level file systems focused on maximizing data throughput.

Flash opens up a range of new possibilities, especially in a VM-aware storage (VAS) infrastructure. New features layered on top of flash can provide guaranteed VM performance and VM-level QoS, real-time actionable VM-level analytics for very granular insights (limited for HDD-based OS systems) and tight integration with the virtualized application and cloud ecosystem for greatest ROI from cloning, thin provisioning and advanced data management.

Tintri: Advancing the 5 core storage rules with VAS

Tintri designed its technology around VAS, leveraging flash benefits, ideal for any type of hardware architecture. Tintri VAS guarantees VM and application performance built around QoS. Historically there have been issues around performance guarantees with QoS set at the volume or LUN level. With Tintri's VAS file system, data from the hypervisor maintains its association with the source VMs, allowing the dedication of separate data I/O paths for each VM. For example, with a system housing virtualized apps such as servers for SQL, VDI and Exchange, VAS can schedule I/Os based on the optimized block size for each application to improve their performance and avoid noisy neighbor/blender effects. While this can be done with brute force creating dedicated LUNs and sub-LUNs, it is tremendously time consuming. VAS makes this process transparent, fast and easy.

Another benefit is root cause analysis (RCA). When an application team complains about VM performance, the storage admin using traditional file systems can only view volumes or LUNs, not VMs. This aggregated view may not show any overall issues because of blending or averaging effects and insufficient view granularity among the VMs it contains. Delving into and solving this could take numerous hours or days to troubleshoot using trial and error with VM placements

Tintri VAS can see end to end of a virtualized environment, supplementing hypervisor latency data to provide detailed, VM-specific visibility that might show to the storage or virtualization admins that all of a VM's latency challenges are coming from the host. This helps quickly identify that there is either a host hardware problem or, as in one example, insufficient allocation of physical memory on the VM. Tintri VAS allows rapid identification of this issue, demonstrating the need to update the host physical memory configuration for that VM. AFAs and VAS provide the level of granularity to rapidly perform RCA and implement remediation.

Tintri's form factor for delivering all these benefits for virtual environments is an AFA appliance, containing both controllers and storage media onboard with always-on inline data reduction services enhanced by integration with leading hypervisor and cloud managers such as vSphere, vRO, SCVMM and Openstack. With Tintri, scaling up is simplified by adding more arrays, avoiding potential performance degradation from adding shelves to an existing controller. And by leveraging commodity hardware supporting our innovative software and VAS, Tintri can offer very aggressive pricing for all levels of capacity.

FEATURE	BENEFIT	
	VM-aware Storage (VAS)	<p>Guaranteed VM performance and VM-level QoS</p> <p>Granular insights from real-time, actionable VM-level analytics</p> <p>Tight integration with the virtualized application and cloud ecosystem for greatest ROI from cloning, thin provisioning, and advanced data management</p>
	All-Flash and Hybrid options	Ability to optimize storage based on infrastructure configuration
	VM as base unit of management	<p>Never deal with LUNs, volumes, or other remnants of outdated storage for dramatically simplified storage management</p> <p>Simplified and accelerated root cause analysis</p>
	VAS architecture	Scale from TB to PB with one employee, set QoS thresholds for critical VMs, run multiple concurrent hypervisors on one storage platform for lower cost
	Appliance-based solution	Scale with no storage controller bottlenecks and minimum latency
	Dedicated data I/O paths for each VM	Eliminate contention and boost performance 6x
	17 TB to 10 PB scale out	Intelligent architecture delivers seamless simplicity and performance with minimal administrative time

Conclusion

Like virtualization before it, flash is remaking how IT infrastructures are architected and managed. The accelerating performance, cost and operational benefits are on a path to make it the dominant storage medium in the near future. But like any technology, integrating it successfully into virtual infrastructures is not a trivial exercise. The choices made on how to do this can have long-lasting repercussions on IT optimization.

The five core storage rules outlined in this paper provide guidance on making these choices to generate the best outcomes. Building a flash strategy around them helps avoid major pitfalls and extra effort while accelerating the realization of performance, TTM and cost benefits. Tintri has built its VAS platform around these five rules for virtual environments, providing a transparent, fast and easy means to tap into all the benefits of flash and the capabilities of a variety of hypervisors. With Tintri VAS, squeezing all the benefits out of a virtual infrastructure is both simple and low-risk.