A SolidFire Insight

Why Is All-Flash Adoption Growing So Fast?

This article was originally published on TechTarget (searchsolidstatestorage.techtarget.com) and authored by George Crump, Lead Analyst at Storage Switzerland.

Flash storage has become the perennial odds-on favorite for IT professionals looking to solve storage performance problems, and there are plenty of flash implementation options available to them. These options include installing flash in servers, mixing flash with hard drives in a shared storage array or all-flash array (AFA). Of these, AFAs seem to be resonating the most with initial adopters. This early lead by AFAs is something that appears counterintuitive given that the hard disk and hybrid alternatives should be less expensive. Another alternative, converged infrastructure, claims to not require a dedicated storage network and hybrid arrays use hard drives. When the challenges that alternatives create are considered, AFAs have earned their popularity by providing a more consistent level of high performance along with easier design and ongoing management.

The Server-Side Flash Challenge

Server-side flash started exclusively as a point solution. An organization facing a performance problem installed a flash board into the server and either through caching software or simply by copying the entire application to it, alleviated the organization's performance issue. When used for just a few servers, server-side flash was a cost-effective method to address an immediate problem but at scale, it became expensive and challenging to manage. Server-side flash was the natural starting point for a flash journey that would typically end with one of the shared flash options.

The server-side flash journey has been interrupted lately by flash-based hyper-converged architectures. These hypervisor-based designs work by aggregating flash installed on the physical hosts to create a virtual flash pool. The benefit is the architectures still can leverage cost-effective server flash and eliminate the need for a dedicated storage network.

There are multiple problems with this approach, beginning with ensuring predictable performance. "Shared everything" infrastructure makes it difficult to deliver that consistency. A spike in application use can cause the allocation of CPU and memory to be out of balance and unavailable to the hyper-converged storage software, directly impacting performance.

Hyper-converged has a role to play. In small data centers the likelihood of a "run" on performance to the degree that it would impact the hyper-converged storage software is rare. In larger enterprises, hyper-convergence will end up converging the data center down to two layers—a compute tier and a storage tier—since it is very rare for these two resources to grow in lockstep.

NetApp®

Hybrid Challenges

Hybrid storage mixes hard disk drives and flash drives into a single chassis. Most of these systems provide intelligence that moves data between the flash and hard drive tiers. The primary challenge with hybrid, again, relates to performance predictability.

Hybrid systems attempt to balance performance and costs by leveraging hard disks in conjunction with flash. The hope is that the majority of data access will come from the flash tier, hiding slow hard disk performance. The problem is that the performance delta between these two forms of storage is substantial. If there is a tier miss, the gap in performance between the two tiers often is often noticeable to users and may lead to complaints.

Some hybrid vendors have tried to address the predictable performance problem by selling extra flash capacity and by offering volume pinning capabilities with their systems. Additional flash capacity reduces the cost advantage of hybrid arrays. Pinning capabilities may reduce the predictability problem, but they also increases flash capacity consumption and require additional administration time, as storage planners fine tune workload placement.

The All-Flash Panacea

AFAs are a performance panacea. An AFA provides a significant performance boost to the data center, helping administrators address performance limitations on even the most demanding applications. By leveraging features like compression and deduplication AFAs make the consolidation of workloads economically comparable to HDD based solutions. AFAs also are popular because they eliminate performance tuning specifically and general storage management as a whole. Every storage administrator can tell war stories of hours and even days spent troubleshooting and fine-tuning storage performance. AFAs effectively eliminate storage tuning and return untold hours to the storage administrator's day. It is enticing administrators to move additional workloads to an AFA.

The Importance of the All-Flash Package

The potential role of an AFA to consolidate all production data onto a single device means it needs to meet a high set of standards. Flash media is the common denominator, but the storage software and hardware that surround it are critical to meeting these high standards. As mentioned earlier, scale-up AFAs are typically purchased to address a specific performance problem but as workloads are added to the array the limitations of this initial purchase become obvious. Scale-out architectures offer some relief to the problem of mixing workloads, but should be combined with intelligent provisioning of the various resources within the architecture to maximize the AFA approach.

Conclusion

Of the available options for implementing flash, AFAs seem to be the favorite of IT professionals. Consistently, the reason given is the simplicity they provide by eliminating performance concerns and the associated tuning that those concerns cause. But it is important that IT professionals look for the complete all-flash package that is both scalable and intelligent in the provisioning of performance so that it meets today's demands while enabling the next-generation enterprise.