

Separating Performance and Capacity Buying Decisions


What do IT managers do when their applications' appetite for storage performance outgrows their backend storage infrastructure's ability to deliver it? Conventionally, the industry approach has been to throw more disks at the problem. This approach is manifested in three different "solutions":

1. Adding more disks
2. Tiering subsystems
3. Adding NAS heads

Adding More Disks

This approach is based on the premise that more spindles means more performance—better yet, more 15K RPM spindles or SSD drives. While this can increase performance of individual subsystems, unfortunately, this approach requires the most expensive drives on the market. Moreover, to realize consistent and predictable performance, IT managers must deploy this solution across all of their backend storage. Clearly acquisition of expensive disks increases CAPEX. Moreover, the addition of higher performance disks drives OPEX increasingly higher as the requirements for space, power and cooling soar.

Adding More Disks

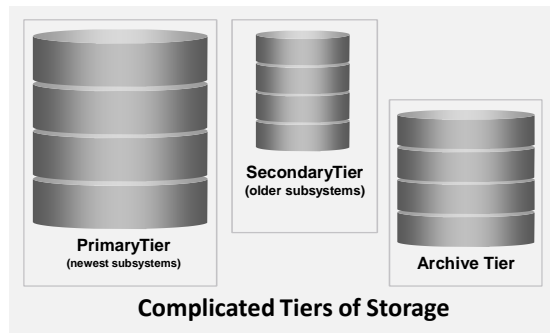


Tiering Subsystems

This method tries to solve the application performance problem through purchasing a faster NAS head with faster disks and logically place it in the network alongside the original "slow" storage. Prior to an application running, IT staff manually copies data from the slow storage into the new fast storage. Sizing the new storage follows the same rules as *Add more disks* – to get to a specific performance level add high cost, high performance disks. This solution does provide


higher performance storage but can support only a very limited number of applications at any one time. Moreover, tiering has no built-in intelligence—this approach has no way of automatically knowing what data is important and when. Coupled together tiering alone has narrow benefits and drives OPEX as more IT staff is required to manually move data back and forth between tiers.

Tiering Subsystems



Adding NAS Heads

Realizing that adding more disks and tiering the data won't deliver the required performance, IT managers' ultimate solution is to keep scaling out the storage infrastructure by adding more NAS heads and their associated disks. Sizing the new NAS heads with storage follows the same rules as tiering the data and adding more disks—to get to a specific performance level per newly installed NAS head, add high cost, high performance disks. Adding NAS heads can increase performance up to a point. However, this is clearly the most CAPEX intensive of the three approaches discussed – purchasing scalable, high performance NAS solutions is expensive. The solution is also very OPEX intensive; migrating data from old storage to the new is time and labor consuming and as with adding more disks; space, power, and cooling add up.

Adding NAS Heads

Each approach has different issues with both OPEX and CAPEX, *all* suffer from a common problem that performance is being increased through the addition of disks. With network storage disk utilization at around 40% (according to Enterprise Strategy Group), IT managers are throwing away 60% of their storage capacity.

Don't buy capacity to get performance


Storspeed's innovative approach to application performance is the SP5000 Application-Aware Cache.

The Storspeed application-aware caching solution is an intelligent, clustered cache appliance that is installed in the network between servers and networked storage. It inspects all traffic between the two and provides an immediate view into storage traffic, application data access, as well as storage performance and bottlenecks. Once access patterns are identified, performance profiles are created to determine exactly which clients, applications, data or file space should be accelerated, thus providing fast, intelligent, and targeted performance improvements. Moreover, as access patterns change so does the accelerated data—automatically, no application, no storage, and no network tuning required.

Sitting in-line between clients and multiple file-based storage subsystems, the Storspeed system intelligently delivers application-defined performance regardless of the backend storage. With the SP5000 Application-Aware Cache, IT managers can deploy additional SP5000s to increase performance without deploying additional disks – which enables IT managers to take advantage of the 60% of unused storage capacity. Storspeed's technology is not only driving disk storage towards full utilization, it is fundamentally changing the IT managers' purchase decisions. IT managers can now deliver enterprise-class performance from solutions typically used for “bulk” data. Private storage clouds and SATA-based storage solutions normally associated with slow

but inexpensive storage can be deployed with Storspeed to deliver the high performance required by today's applications

Consolidated Storage with the SP5000



Storspeed customers outperform today's high end NAS Servers while utilizing existing file-based storage subsystems and commodity components—no fork-lift upgrades, no mount point changes, no data migrations, no configuration changes of any kind. Storspeed has broken the cycle of buying capacity to get performance—for the first time IT managers can purchase the Storspeed system when they need performance and commodity disk solutions when they need capacity.

Deploy the lowest-cost, densest storage you can find

Storage managers can no longer afford to scale their storage performance by buying capacity. The future of storage requires meeting the growing application demand for performance while fully utilizing the capacity offered by the lowest-cost, densest storage available. The only way to take advantage of this model is to scale performance and capacity in an intelligent way. The Storspeed SP5000 is the revolutionary solution that delivers intelligent application performance while allowing IT managers to deploy low-cost, high density commodity storage.

www.storspeed.com