



Why you might not need a new storage array

In today's datacenters, many factors are putting significant pressure on storage systems. Virtualization has gone mainstream, with most organizations having a "virtual first" policy for bringing new applications online. This means more applications are competing for the same resources, and that traditional I/O handling techniques are less effective, with multiple applications' I/O streams combined by the virtualization "blender effect." The throughput and transactional requirements of enterprise applications and databases can be further exacerbated by virtual desktops, whose interactive workloads add increased pressure for low latency delivery of applications from storage.



A range of options

Many organizations start to address lagging storage performance by adding more drives to their storage arrays. But adding shelves of drives can increase complexity, as well as power and cooling costs. With this strategy, eventually storage systems hit a maximum drive count and need to be upgraded.

Another common solution to storage system performance issues is introducing flash technologies like PCI-e cards or SSDs. However, just putting SSDs into existing storage systems can exacerbate the problem, as there is a significant amount of overhead associated with efficiently supporting flash devices. Filling disk arrays with flash can also worsen the problem by driving more IOPS to disk than the processors in the array were designed to handle.

Alternatively, some organizations choose to put flash devices into the specific

servers that have performance-sensitive applications. While this approach can address known performance bottlenecks, it often requires significant manual configuration (like placing certain database tables on a particular device). It also isolates performance benefits to a specific part of the environment.

Often this is the point at which IT organizations look at replacing their storage array with something more powerful. All-flash arrays advertise supporting millions of IOPS, and hybrid arrays are positioned as an affordable interim choice. However, the expense of either of these solutions can be compounded by the operational costs of a new array: data migration, learning new snapshot and replication tools, rewriting scripts and reports, and retraining staff.

A lightweight solution: Storage acceleration

There's another option for organizations eager to improve performance in storage systems without the cost and complexities of flash or the expense and disruption of a new array. While storage systems have been buckling under the increased pressure of additional applications, increasingly more powerful servers have evolved to support these applications. The amount of addressable memory in today's CPUs is 3X what it was just five years ago. And the number of cores available in each CPU socket is growing exponentially every few years. Similarly, 10GbE Ethernet is mainstream now, with 40GbE close on its heels.

Server-side storage acceleration solutions enable IT administrators to leverage the memory, CPU, and networking resources – abundant on application servers – to take some of the burden off of storage systems. Here's why its adoption rate is growing so quickly:

1. Resources are available at commodity prices when they are in a server. The cost to end users of adding cache or processing capability to a storage array is much higher than the cost of adding the same resources on the server side.
2. Organizations can preserve their investment in existing systems. Many storage acceleration solutions don't require changes to storage configuration, hypervisor configuration, scripts, operations, or other processes.
3. When I/O is processed on the server side, latency is significantly lower. Applications are closer to the media from which they are requesting data, which means they can get responses much more quickly.

With storage acceleration, a significant portion of I/O requests never reach the storage system; instead they are processed on the host servers. When it is I/O performance (not storage capacity) driving the purchase of a new array, storage acceleration enables organizations to address storage performance concerns without purchasing unnecessary capacity.

Preserving investment in existing storage systems

Leveraging storage acceleration enables organizations to continue using their existing storage systems while gaining the added benefit of increased performance. Some organizations find that they can keep their storage systems for longer, delaying an outlay of capital. They are also able to avoid the costs associated with data migration and the operational changes that come with a new array. For other organizations looking to implement emerging storage technologies, storage acceleration can enable IT to "buy time" while formulating a bigger strategy around datacenter architecture.

Another key benefit is the ability to support more workloads on the same infrastructure. This can lead to a number of advantages. For some, it's an increase in the number of applications they can support, which can speed virtualization adoption or support business expansion plans. For others, it's an increase in the number of users they can support, which can speed VDI adoption.

For all IT organizations taking advantage of storage acceleration, the result is more flexibility and less upheaval. While storage acceleration solutions may vary in their implementation details (for example, scalability, their ability to leverage existing infrastructure, and how disruptive they are to implement) they provide great value to organizations looking to get more from their existing storage systems.

For a storage acceleration solution that can be installed in under 30 minutes and doesn't require any new hardware, visit www.infinio.com