



Exploring a Scale-out Storage Solution

Storage is hotter than ever before, and there is no shortage of options for organizations looking to boost their capabilities. Find out how to choose the right storage solution and the benefits proper storage technology can produce.

By [Derek Schauland](#)

Storage is a hot topic in technology today, with all-flash arrays and spinning drives, cloud-based architectures and hybrid storage. In short, there is no shortage of ways to store information. The biggest consideration isn't just choosing a disk solution, as all of the providers can store data. It's understanding what is unique about each option. Organizations must consider the additional value that storage vendors can provide. The unique value the providers can offer that sets them apart is where the prime consideration will be for customers, not in the data storage portion of the product.

Choosing the right storage solution which will provide the best fit for your organization requires consideration of several elements:

The unique value the providers can offer is where the prime consideration will be for customers.

Workloads: What things will the storage be used to accomplish? What are you going to do with it? Will it be used handle files for the organization or stand up virtualized workloads?

I/O: Once workloads are identified, what is a tolerable I/O level for the organization? Keeping the storage at optimum performance to accommodate workloads and user requests is paramount. Remember, the storage is only as fast as the users accessing a given workload think that it is.

Scalability: Can the storage grow faster than you need it to? While it may seem a bit overzealous to plan for storage to grow faster than is needed, but being able to scale at the drop of a hat can end up being very important. Scalability is more than just adding spindles to grow the storage, it also requires adding more nodes to allow for maintenance and updates without affecting (or limiting effects) to the workloads.

How can the storage provider add value to the primary considerations of storage?

Major Trends

Cloud Storage

The cloud has become the new disk drive in many cases. Because it can be accessed from almost anywhere, the convenience and productivity that is allowed when considering the cloud as a storage and application platform is almost unheard of in previous environments. Not only does

the cloud enlarge and organizations office, but it can also reduce maintenance overhead that may be associated with backup and disaster recovery as well as general power and cooling considerations. Building a storage solution to rapidly scale and accommodate continuous workload additions keeps user experience high at all times.

Virtualization

Considering the needs of almost all organizations to do more with less, especially in terms of working in the cloud, virtualization helps maximize utilization across fewer resources. But with increased utilization can come increased pressure on storage. Remember that the shared storage that backs these solutions often has an increased workload because the number of virtualized applications increases with better utilization while storage resources remain static. Features like deduplication can increase the load on back end storage as well.

Storage may be hot, but if you look at the hype, flash is hotter.

What about flash?

Storage may be hot, but if you look at the hype, flash is hotter. The speed of the disk in an SSD is pretty amazing. For many customers, the first experience they have with SSD is in their laptops, where things like boot time alone certainly impress. When you consider that technology in enterprise storage, the possibilities increase considerably. The challenge with flash is deploying it in a cost-effective way that best accounts for its architectural differences from spinning drives.

New Architectures

There are several new architectures emerging in the datacenter that support and leverage these trends. New storage systems (either all-flash or hybrid arrays) leverage flash on the backend to optimize application experience. In hybrid arrays, flash often assumes the role of cache to help keep the active working set on the fastest media possible. In all-flash arrays, flash is also used as the storage medium.

Other architectures, hyper-converged systems for example, bring storage, networking, and compute into the same enclosure. This allows both the functionality and management of the system to be consolidated into a simplified interface and plays well in IT shops where specialization in each piece of the stack is difficult. The adoption of these technologies is high in remote office and branch office locations.

Another emerging architecture is server side processing, which can speed performance and improve perception by those using the system by allowing the back end storage to process less data. This can also improve the overall life of backend storage and improve performance. These solutions are an alternative to expensive arrays-side flash solutions, giving IT decision makers the option to avoid jumping right into a flash configuration if it isn't absolutely needed.

What does Infinio bring to the enterprise storage party?

Infinio offers a server-side processing model for storage without the addition of any hardware into an environment. It optimizes data transfer between hosts and storage using caching to improve performance. Because the data processing occurs across host nodes within an environment, the I/O load placed on the traditional storage array is reduced.

Consider for a moment how much processing takes place on a traditional storage system, and how much of that processing is of duplicate data. With Infinio, repeated read requests of the same content are served out of local cache rather than off the storage system. This reduces latency as well as network load, and results in less overall processing on the storage array.

In short, Infinio can offer the performance advantages of all-flash arrays without the disruption or expense of implementing a new platform. By separating storage performance from storage functions like capacity and data protection, Infinio enables IT organizations to manage performance as an independent resource. Read on to find out how Infinio accomplishes this without adding any additional hardware to an environment.

How does this work?

Infinio uses available RAM across nodes in a cluster to process data in a distributed fashion. Each node in an Infinio cluster provides a small amount of RAM, which is then aggregated into a global, deduplicated cache across nodes. Each 4K block of content in the cache is only stored once, and is accessible to all of the virtual machines on all of the nodes in the cluster. This is an effective architecture with just small amounts of RAM because of the impact of deduplication.

Infinio can offer the performance advantages of all-flash arrays without the disruption or expense of implementing a new platform.

“Infinio is making use of memory, which is in the nanosecond range and is becoming more prevalent across servers – we’re taking advantage of these advances in networking where remote memory access can actually be equal to or faster than flash.”

Scott Davis - Infinio

When applications request data, if it's in the cache, it is returned very quickly.

When applications request data, if it's in the cache, it is returned very quickly, without a call to the storage system. If it's not in the cache, then data is returned from the storage system like during normal operations. The net effect of this is that cache hits are returned very quickly, as they are local to the server cluster, while overall storage performance on the storage increases, as a significant number of requests are offloaded from the total workload.

In addition, the cache created by Infinio is completely transparent to the systems using the cache. Typical operations such as backups, snapshots, replication, and VAAI integration remain unchanged. Administrators don't need to learn any new systems or tools, and existing scripts and reports are still operational and valid.

What are the best use cases for this type of technology?

The best examples are situations where much of the content is the same across many workloads. Infinio's deduplication allows for more efficient caching using small amounts of physical space. Virtual Desktop Infrastructure (VDI) is a great example of a use case, since files like notepad.exe are generally the same across each VDI session, and once they are stored in cache, they can be accessed by any virtual machine on any host. While Notepad.exe was used as an example, many files and parts of files across Windows images or VDI sessions are the same, which is where the deduplication really shines to improve the experience.

A better example of deduplication might come from a PowerPoint deck that is shared by an organization. The template for the deck, containing the logos and general color scheme for the slides will be the same in many cases, with only graphics and things pertinent to a given presentation changing. Since Infinio works on a block by block basis to deduplicate content in cache, the original template and pieces used across many PowerPoint files will be deduplicated and stored only once. This way, more content can be stored in a cache with an otherwise small physical size.

Caching and storage optimization provided by Infinio can reduce the overall cost of optimized storage and perhaps hold off the transition to flash for many workloads. [VR](#)

Derek Schauland has worked in technology for 15 years in everything from a help desk role to Windows systems administration. He has also worked as a freelance writer for the past 10 years. He can be reached at derek@derekschauland.com.

To learn more about Infinio and server-side storage performance, go to www.infinio.com/solutions

Commissioned by
INFINIO