



SWIFTSTACK WHITEPAPER

Top 5 Things You Should Know About
Unified Storage

November 2014

The increasingly heterogeneous makeup of storage infrastructure has made life difficult for storage administrators and data center architects. Many organizations today are supporting both network-attached storage (file-based) and storage-area networks (block-based), but usually for different workloads and applications.

As a result, these organizations are struggling to achieve essential goals such as interoperability, cost efficiency and long-term scalability. Add in the confusion with the introduction of object storage as a way to help deal with ever-growing mountains of unstructured data, and it's easy to understand why IT organizations need new solutions to address different storage architectures in a comprehensive and coherent manner.

This is where unified storage comes into play. Most IT professionals know that unified storage architecture allows them to run and manage files and applications from a single storage system, and they are becoming more aware of and interested in integrating object storage into their storage architecture strategy. This is particularly significant in cases where IT organizations have stepped up their use of virtualization or data center infrastructure consolidation initiatives.

The benefits of such a solution can be very impressive. For instance, legacy storage environments have often evolved through multiple, disparate purchases of either network-attached storage (NAS) or storage area network (SAN) solutions. By comparison, moving to a single unified storage solution can save considerable money on capital expenditures. It also can dramatically simplify storage management when compared with having to manage multiple systems based on either file or block architecture. Today's unified storage solutions also support popular features such as snapshots, data deduplication and replication.

But the attraction of such a solution means that some storage vendors are claiming to have fully unified storage offerings, but they don't offer a truly integrated and optimized approach to storage unification. Since the concept of unified storage has great appeal to storage professionals, it's natural that suppliers want to position their products in that way—even if they don't all achieve unification in the same way.

What to Keep in Mind

Here are five things IT professionals and storage administrators should take into consideration about unified storage when looking for new ways to achieve consolidated architectures for cost efficiency and simplifying storage management:

- 1. Not all unified storage solutions provide full integration of file, block and object storage.** Having a "unified" storage solution sounds like—and certainly is—a good thing. So it's not surprising that most storage vendors are marketing their solutions as being unified. But buyer beware: Many "unified" storage systems are block-based storage systems tweaked or retrofitted to also provide file access across Network File System (NFS) or Common Internet File System (CIFS). This approach fails to optimize the unique performance requirements of different workloads. A SAN that can also provide file-level access is a less-than-efficient approach for handling file access, while file-based NAS providing block-based access is unlikely to deliver sufficient and sustained performance for IO-intensive workloads.

2. **Don't assume that results from proof-of-concept deployments will hold true as you scale to production environments.** While lab-based testing of pilot deployments are helpful first steps, don't be lulled into a false sense of security. It's very difficult to simulate real-world environments made up of unexpected spikes in user requests, fluctuating access demands for performance-intensive use cases, network bandwidth limitations or unexpected availability problems. In the less-taxing environments characterized by proof-of-concept requirements, some unified storage systems may perform fine but then suffer availability or latency problems when conditions change, such as when individual workgroups spin up new virtual machines (VMs).
3. **Support for object storage significantly changes the rules of the game in unified storage.** Object storage is an exciting and rapidly growing segment of the storage landscape because of its ability to handle fast-growing volumes of unstructured data. Object storage's near-limitless capacity expansion, combined with its ability to sustain IO performance as requirements scale out, make it an increasingly important part of a data center's storage roadmap. In unified storage solutions not designed or optimized for true file, block and object storage integration on a single platform, organizations will probably not see many of the full benefits of object storage, such as reducing capital expenditures and simplifying storage management.
4. **Ask potential suppliers how they overcome the inevitable performance latency problems.** One of the most attractive benefits of a truly unified storage solution is its ability to support different workloads on a single platform—even if some of those workloads have very different performance requirements. There are some workloads where delayed access to files, blocks or objects can literally cost organizations huge sums of money, such as online transaction processing. Storage vendors whose solutions are really dedicated NAS solutions, with the ability to provide block-level access, will undoubtedly run into situations where enterprise workloads will not perform nearly as well as they would in a system optimized for use cases like databases or enterprise-wide virtualization.
5. **Virtualization—especially VDI—severely taxes storage solutions that aren't designed from the ground up for unified storage.** As more organizations embrace virtualization to reduce complexity, trim capital expenditures and improve IT flexibility, unified storage solutions can be a great asset. Unified storage—truly unified solutions, that is—aligns nicely with virtualization's core principle of allocating IT pools like storage to be accessed as needed. When done right, unified storage provides organizations with the ability to select how to store VM-based data without having to provision dedicated SAN and NAS devices for unique workloads. However, IT managers and storage administrators need to keep in mind that storage pools must scale not only on the basis of capacity, but also especially on performance required of many VM-based workloads. It's not unusual for IT departments to discover that all those VMs that have been spun up—often without their knowledge—can severely tax the performance limits of storage systems not optimized to handle enterprise-class virtual workloads on a single storage system. This contention for storage resources severely affects system availability and performance.

The SwiftStack Approach

IT decision makers need to keep their options open when it comes to unified storage, because rarely will a single unified solution solve all storage problems and address all use cases. Inevitably, there will be a need to mix and match unified storage with legacy storage systems that may be running unique or lower-tier workloads.

But organizations looking for a flexible, efficient and optimized approach to unified storage should look for a solution that offers such functionality as replication, flash storage support and hybrid workload support. Ideally, this solution would be based on open standards in order to avoid vendor lock-in and give customers the ability to run a unified solution on affordable, industry-standard storage hardware.

One such approach is offered by SwiftStack, which is becoming a leader in software-defined storage. SwiftStack's solution is based on the popular OpenStack Swift object storage engine, which simplifies integration, scalability, management and storage unification even as storage capacity and performance demands escalate.

No matter the circumstances influencing an organization's storage architecture and performance requirements, IT departments need a mix of both file/object storage and SAN (block storage) systems. There is no single solution that is great at solving all use cases, especially at scale. Storage systems that combine object and file storage can deliver great performance and scalability because these two systems are designed for unstructured data. Further, Object Storage is built to scale to petabytes and beyond for a single cluster. Customers can scale I/O throughput and storage capacity individually with SwiftStack, which allows them to scale what they need, when they need it, to address each issue.

SwiftStack's object storage solution was purpose-built for storing unstructured data in either file or object formats, as well as delivering superior performance in highly scaled environments. This is an important issue to consider, since customers now demand the ability to scale both I/O performance and capacity—not just one or the other.

By separating the physical storage from the data and metadata management layer, the SwiftStack platform enables administrators to manage their object storage infrastructure through the company's SwiftStack Controller dashboard. This provides a single view across all devices and nodes, and enables rolling, nondisruptive upgrades and automated device monitoring.

SwiftStack's approach represents a unified storage solution for storing unstructured data because it is engineered for multiple data types and access methods (CIFS/NFS and Object). As a result, it allows storage resources to be deployed across a single platform that integrates legacy file systems and object-based data storage.

Conclusion

Demands on storage administrators have never been greater. As data growth—especially unstructured data—mounts by huge volumes each year, there is increased pressure to reduce storage costs while still supporting higher capacities and escalating performance requirements in IO-intensive workloads.

The most significant gating factor for storage administrators looking to create unified solutions is the importance of delivering high performance at scale. Combining object storage with traditional NAS/ file-based storage is a great option because this approach to unified storage is designed specifically for unstructured data, which now dominates the data storage landscape and is eating up huge amounts of capacity and is taxing storage systems' performance at scale.

SwiftStack is delivering a solid solution to store unstructured data that can scale both capacity and I/O as workload characteristics require.

SwiftStack's unified storage solution has been established as a viable, reliable and efficient way to deliver on the promise of unified storage. Its support of the increasingly important OpenStack Swift object storage engine is a key differentiator from other legacy storage systems that have been tweaked to deliver incremental support for either file or block storage. As a result, SwiftStack's solution has been adopted by some of the most demanding storage customers around the globe, including eBay and HP.

For more information, [go to www.swiftstack.com](http://www.swiftstack.com).