

Flash Strategy 2.0: Optimizing NetApp with Flash and Cloud



by George Crump, Lead Analyst

For most organizations, databases are the heart of the data center. Not surprisingly, its IT infrastructure revolves around those databases. But there are market segments like Media and Entertainment, Life Sciences, Financial Services and Technology where unstructured data is at least as critical, if not more so, than block-based data. Most of these organizations count on a network attached storage (NAS) system to provide the file services for this unstructured data. The go-to product for these organizations has been NetApp's FAS family of systems. But as these environments have scaled, NetApp's operating environment, Data ONTAP, has struggled to meet performance and capacity demands. To overcome these limitations file centric companies need to purchase additional NetApp filers which increases cost and complexity.

NetApp tried to address customer concerns by integrating flash storage into its hardware and deliver a clustered version of Data ONTAP. While these developments helped, flash integration has been slow to come to fruition and performance on flash-based filers is not as optimal as it could be. Clustered Data ONTAP does bring a sense of capacity scale but also adds latency which adversely affects performance. It also adds complexity.

Leveraging the Cloud to Abstract Performance and Capacity

Organizations that have unstructured data as a critical aspect of their businesses are looking for an alternative to staying in the NetApp eco-system. The problem is most solutions force IT managers to abandon what NetApp does well, namely data services like volume management, snapshots, cloning and replication. Not only does abandoning NetApp require the re-learning of other services and a potential downgrade in capabilities, it also means changing the way back-end processes like data protection work, since many of these organizations have integrated NetApp features into those processes too. In short, many organizations are heavily committed to NetApp and a change to a new platform would be painful.

The answer is to abstract performance from active storage, and active storage from archive storage. Abstraction allows for the optimization of individual components for their purpose; a performance tier accelerates I/O leveraging flash storage to its fullest

capability, an active storage tier continues to leverage NetApp data services and either the public or a private cloud is suitable as a long-term data archive for cold data.

The Performance Tier - Filer Flash Done Right

Flash storage represents a significant change in the way a storage system delivers data to the requesting user or application. In one motion, the slowest part of the storage system, the storage media, is now the fastest. More important than the raw IOPS of flash though is its responsiveness. These devices have almost no latency. The problem is that systems like NetApp, which provides the features the enterprise counts on, has to do too much and actually impedes the potential performance of the flash media. For the first time the storage media is actually waiting on the storage software.

Abstracting performance, both software and hardware, from the traditional storage system is the best way to go. The performance tier vendor can focus their software on extracting maximum performance from the flash media not on adding features that already existing in the active tier. The performance tier complements, not replaces, the active tier. As a result the performance tier can streamline its feature set to further reduce latency and take full advantage of flash storage media.

The performance tier is typically a clusterable flash appliance and can provide performance acceleration to all the organization's NAS systems. Some of these solutions provide a global file system that spans across all the devices, from each tier, both inside the data center and the cloud. The global file system allows all the NAS, both on-premises and in the cloud, to act as one.

The performance tier also acts as a shock absorber to the active storage tier. All inbound write I/O lands on this tier first and then is copied to the active storage tier later. At the same time most recently read data is copied up to the performance tier. The result is that applications and users experience not only the performance of flash storage but the *fully optimized* performance of flash storage.

The performance tier represents a surgical strike against performance problems instead of a wholesale replacement like an all-flash array is. Instead of replacing the active tier it takes full advantage of it.

The Active Tier - Leveraging Existing Investments in Hardware and Software

The active tier is typically the existing NetApp NAS systems. Their primary role shifts from *servicing* data to *storing and protecting* data. All data from the performance tier is also stored on the active tier. The active tier continues to use all the existing snapshot and replication software. Data centers don't need to redesign their backup and other processes.

The combination of a performance tier with an active tier also drives down the cost per GB of the active tier. Without a performance tier, most NAS systems purposely limit

expansions to only 50 percent of their potential capacity. With a performance tier in place the active tier no longer has the responsibility of delivering top performance. As a result it is possible to expand the active tier to its full capacity potential and the highest capacity hard drives can be used.

The Cloud Tier - Driving down the Hard and Operational costs around Long Term Data Storage

The third data type is cold data. Data the enterprise needs to keep "just-in-case" it needs to access that data in the future. NAS systems are overkill for this archival use case. As a result many data centers are looking at cloud storage to drive down the cost of long term data preservation. The performance tier facilitates either a public or private cloud solution. The cloud tier becomes part of the performance tier's global file system so that movement between tiers is seamlessly performed.

Conclusion

When the data center presents most vendors with a performance problem, the vendor's typical answer today is to replace the current storage system with an all-flash equivalent. In the case of NetApp that is one of their all-flash filers. In the case of a potential alternative vendor to NetApp, the replacement is a hardware solution designed to leverage flash. These "solutions" are essentially replacements for what is already there. Abstracting storage by data type (performance, active and cold) and then providing intelligence to automatically move data between storage systems as its use changes, allows the data center to leverage their existing storage investment. It also provides a path to the cloud.

Sponsored By [Avere Systems](#)

About Avere Systems

Avere gives organizations the ability to put an end to the rising cost and complexity of data storage, the freedom to store files anywhere – in the cloud or on premises – all without sacrificing the performance, availability or security of their data.