

# Investing Strategically in All Flash Arrays

*NetApp's All Flash FAS Storage Systems*

**By John Webster**

**May 2015**



**Evaluator Group**

*Enabling you to make the best technology decisions*



**Evaluator Group**

## Executive Summary

The transition from dependence on rotating disk to integrating solid state storage within the enterprise data center has wide ranging implications for enterprise IT. As this transition occurs and the cost per unit of solid state storage capacity inevitably decreases, enterprise IT will realize that solid state storage technology benefits many aspects of enterprise IT. The true value of flash storage is not determined on a cost per TB calculation alone.

Indeed, all enterprise applications can benefit from integrating with flash storage. But in particular, those that are OLTP-oriented as well as those hosted on virtualized servers will see accelerated levels of performance resulting in an immediate return on investment (ROI) as measured by increased levels of application user productivity. The simple fact is that more transactions completed in a given period of time yield more revenue. An investment in solid state storage drives increased levels of application performance resulting in more revenue generating transactions every business day and yielding an immediate return on the flash investment.

The total cost of ownership (TCO) for solid-state storage vs disk is also reduced via lower operational costs. Enterprise IT administrators typically run disk storage arrays at 35-50%, resulting in operational and budgetary inefficiencies. Multiples of these are now being replaced by single arrays composed entirely of flash storage media—the all-flash array—that offer performance gains coupled with gains in efficiency.

An investment in flash storage should result in much more than accelerating enterprise applications; it should at least preserve if not improve upon existing data protection and business continuance processes. It should also demonstrate an ability to be scaled upward in capacity and performance without disruption to application users. A platform choice in increasing demand from enterprise storage administrators is the all- flash array because it can converge all of these capabilities onto a single platform managed through a single management application.

Here we detail the broader, positive impact that SSD and, more specifically the all flash storage array, will have on enterprise IT. We also expose both the business and IT metrics needed for an evaluation of this new storage platform. Last, we review NetApp's implementation of an All Flash Array based on Clustered Data ONTAP and examine its value in light of making a strategic investment in all flash storage arrays.

## The Business Impact of the All Flash Array

The use of flash as primary storage integrated with critical applications and virtualized server environments accelerates I/O performance. For example, with transaction-oriented database applications, an investment in a production data center-grade all-flash array significantly lowers application latency, allowing business users to work more productively, while preserving the required resiliency and data integrity features. This gain in productivity translates directly to an immediate ROI that can be used to justify the investment in an all-flash array.

A positive ROI can also be seen on the customer side of the equation. For business user groups and corporate executives wanting to engage with customers via web- and cloud-based applications for example, flash integration underpins an ability to deliver to these customers a positive application experience, thereby enhancing revenue generation.

However, we believe that there will be more far-reaching impacts on enterprise computing once the broader implications of flash storage deployment are fully understood. These implications include:

### **Increased Business Application User Demand for Performance**

The understanding that SSD accelerates computing performance is becoming more generalized. For example, anyone who buys a new laptop will see SSD as an option and will learn that performance is what it's all about. It is also likely that once a user owns a laptop with SSD, going back to rotating disk is no longer an option. Performance is addictive. A similar scenario will play out among transaction-oriented application user groups. If the performance of some applications can be accelerated, why not accelerate more applications? And if my coworkers get the power, why can't I tap into the power source as well? Once these user groups see an ability to accelerate their work lives, they will never willingly go back.

Storage administrators have historically allocated database indexes and log files to fast drives with RAID 1 for high performance database applications (MS SQL Server 2014, Oracle 12 and RAC, MySQL). Now that flash storage is cost-effective at scale, the entire database can be hosted on an all-flash array for even greater performance and ease of management by consolidating onto one high performance storage tier. The investment in such an array can empower all database and transactional systems users to become more productive. As mentioned, this gain in productivity represents an immediate return on the investment in an all-flash array.

### **Expansion of Virtualized Server Environments and VDI Deployments**

Virtualized server environments make efficient use of CPU and RAM resources but generate random I/O patterns that can have a negative impact on data throughput from disk—a result often referred to as the I/O blender effect. Virtualizing desktops (VDI) on top of virtualized servers only exacerbates the problem; the all-flash array neutralizes this effect. All VDI users can be given at least the performance level they were used to with physical desktops if not better. More importantly, consistent and predictable performance levels can be maintained—a critical issue with VDI deployments—as these environments grow. Upward scale can be achieved without adding more server hardware, allowing the VDI environment to grow economically and with minimal impact to existing application users.

Virtualized servers also pose challenges for IT administrators responsible for data protection and business continuance. The all-flash array that is optimized for data resiliency and availability within virtualized server environments can not only address these issues but potentially accelerates the performance of these processes.

## Keeping Web-Hosted Application Users Engaged

The challenge of keeping web-hosted application users engaged is of particular importance as the ability to use mobile devices in the context of enterprise business functions is now in demand. Deploying the all-flash array for web-facing and mobile applications assures performance for users under highly variable workload conditions. An investment in an all-flash array allows enterprise IT to take on additional users for revenue growth while scaling performance to meet new demands. When a website is highly responsive, users are more likely to stay on that site and can be kept from straying over to a competitor's site.

In all of the above cases, the investment in an all-flash array can lead to better business-related outcomes, driving more revenue per day, week or month; attracting and maintaining loyal customers; and making application users more productive.

## Economic Benefits to IT Resulting from Consolidation, Greater Server Efficiency and Storage Array Longevity

At the moment, enterprise IT is generally deploying all flash arrays to accelerate and/or stabilize specific application workloads such as database transaction processing (OLTP) and virtualized server deployments including virtual desktops. And as enterprises are now taking increasing advantage of converged and hyper-converged systems, all-flash arrays are become the storage system of choice in these deployments as well. The drivers of this trend commonly include:

- The ability of such an array to take maximum advantage of multicore processor-based servers
- Consolidation of multiple high performance disk arrays to a single flash-based array
- Elongating the time between platform refresh cycles that are often costly and disruptive—a change from three to as much as five years between array replacements that is enabled by continuing improvements in flash module longevity

## Operational Budget Savings from Enhanced Efficiency

The proliferation of high performance, multicore servers has created an imbalance characterized by an overabundance of processing power on the one hand with an inability to fully utilize it on the other. The all-flash array introduces a persistent storage option that is an order of magnitude less expensive than DRAM and at least an order of magnitude faster than disk. By eliminating rotating disk as an I/O bottleneck, server-based CPUs spend less time waiting for data and more time doing actual work; thereby putting the imbalance between CPU power and actual output back into balance. This means that servers can handle additional workloads without adding CPU performance yielding on-going savings in application support costs.

More importantly though, because many software vendors are now basing licensing costs on the number of CPU cores in use, getting more work out of each individual processing core saves

considerably on expenses devoted to software licensing fees as the workload grows. These savings, compounded over time, represent a significant return to the IT budget on an all-flash array investment.

Finally, operational expenses unique to database environments can be reduced. Database applications typically generate a random I/O workload. This creates a problem for disk arrays that is solved by increasing the number of spinning media devices for the aggregated number of I/Os. Using flash instead eliminates the need to continually add disk spindles to an array to maintain consistent and predictable levels of low-latency performance as the database scales to meet increasing workload demands.

## Data Center Storage Array and Server Consolidation

While it is not always the case, additional returns can be realized when it becomes possible to consolidate the workloads of multiple rotating disk arrays to an all-flash array. The ability to consolidate will be determined by application performance requirements and the degree to which rotating disk arrays that are candidates for replacement are being utilized. For example, arrays that are only loaded to 30% of total capacity for performance reasons are prime candidates. If even two of these are consolidated to a single all-flash array, the software licensing and support costs for one of these arrays is eliminated—more if others are included.

The fact that enterprise IT will likely be migrating desktops from Windows XP to more recent versions of MS Windows and that Windows Server 2003 is reaching end-of-life will likely require a large scale replacement of servers to handle the increased load. Consolidating storage to an all-flash array for these projects could result in users seeing an increase in performance and allow administrators to deploy fewer servers than might otherwise be required by making more efficient use of CPU resources as discussed above.

## Extending Useful Life

One of the points of deploying solid state storage is to replace inherently unreliable mechanical devices full of moving parts with solid state devices that are without them. With continuing improvements in flash module durability, IT administrators can now extend the useful life of an array from three to as much as five years. Doing so reduces over time the data migration and application downtime costs associated with array refresh cycles. And when clustering technology can be applied to the all-flash array platform, even these costs can be reduced if not eliminated.

Replacing mechanical disk with solid state offers additional opportunities for cost savings that can be used to offset the investment in flash storage:

### Power/Cooling

From a power cost perspective, flash can save significantly over an extended time period vs spinning disk and can reduce the amount of power required by the data center as it is scaled to meet increasing processing demands. Over several years, these savings contribute to a real reduction in ongoing operating expenses compared to what would be expected with hard drives.

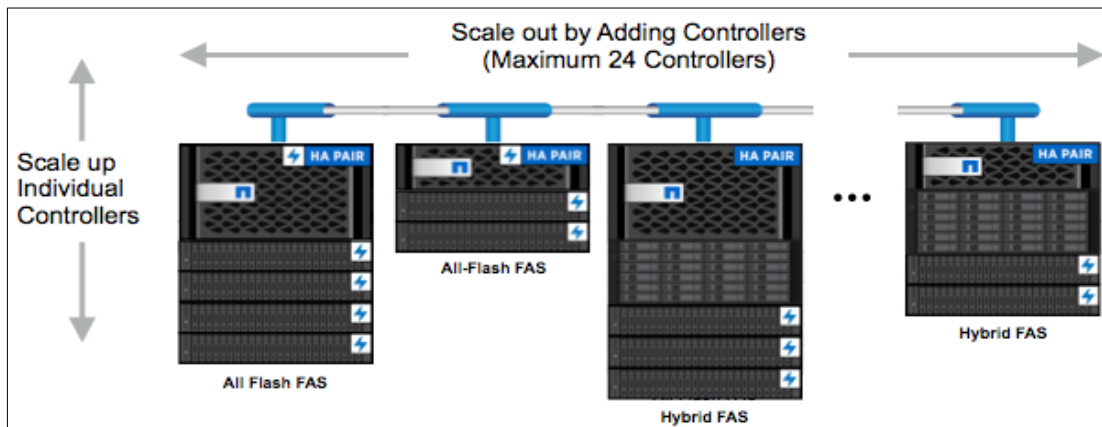
## Lower Maintenance

Flash solutions also contribute to lower operating costs by reducing the maintenance required to periodically replace spinning disks. As more disks fail, there are more disks to replace. This ongoing maintenance drives up the TCO of traditional hard disks compared to flash.

## NetApp All Flash FAS

Given the increasing demand for all-flash storage arrays, a number of vendors now offer them for use in production data center environments. However, rather than introduce an all-flash array with a different operating system and software-based features vs. their existing array product lines, NetApp has chosen to base their All Flash FAS array on its existing and widely deployed Clustered Data ONTAP platform. As a result, customers familiar with Clustered Data ONTAP are not required to learn a new storage operational and support environment when transitioning from rotating disk to flash.

The NetApp All Flash FAS portfolio currently consists of four model variations with flash storage capacities from 96 to 384 TB. All are designed to be integrated with a clustered, modular hardware architecture running clustered Data ONTAP and can be deployed in both SAN and NAS environments in high availability (HA) pair configurations. All Flash FAS clusters can be scaled up to 24 nodes for a total of almost 5 PB of flash storage capacity. In addition, a single cluster can be built on all-flash storage technology or can be composed of a combination of flash and disk arrays in a hybrid FAS configuration (see Figure 1 below).



**Figure 1:** A mixed cluster of All Flash FAS and hybrid FAS nodes. The clustered Data ONTAP global namespace functions across all nodes. (Source: NetApp)

Multicore CPUs and large system memories in FAS storage controllers are designed to take maximum advantage of flash performance. From a hardware perspective, up to 40 CPU cores and larger memory sizes are available per node to manage data in order to assure a balance between controllers and flash modules for I/O processing thereby reducing “hot spots” and extending the life of flash devices.

In adding All Flash to the FAS system series, NetApp has preserved the FAS bi-dimensional scaling architecture—the ability to scale storage systems out (add nodes to a cluster) and up (move to the next platform up within the same family) at the same time<sup>1</sup>. Bi-dimensional scaling means that performance across the cluster can be maintained while demands for increased storage capacity are met. The Clustered Data ONTAP global namespace that functions across all nodes—flash *and* disk storage nodes—is maintained without interruption and non-disruptive operation is assured as the system is scaled. In addition, Clustered Data ONTAP will redistribute active data across flash or disk nodes as they are added or replaced based on performance needs.

## Scale-out Cluster Advantages

An important capability included with All Flash FAS is the ability to scale-out the cluster non-disruptively by adding flash-based cluster nodes. Customers can start with deploying flash nodes to deliver high performance for a dedicated workload for example. If additional performance is required, customers can scale-out a cluster—up to 24 nodes with up to 384TB of all flash capacity per HA pair. When additional capacity that doesn't need to be of the highest performance is also required, disk arrays can be non-disruptively added as well. This type of storage clustering combines high-performance flash nodes with disk that can become a single storage repository for multiple workloads—high performance and high capacity—sharing the same data protection and business continuance capabilities and under a single application for operational management.

FAS clustering also enables non-disruptive movement of workloads to the node that best meets the application user's performance expectations while preserving IT's need to achieve price/performance efficiency. IT administrators can non-disruptively update the entire cluster in a modular fashion over time by retiring nodes from service and replacing them with faster, denser nodes while data is transparently migrated to other nodes to maintain normal operation, and then rebalance data across the cluster once the new node is on line. A significant result will be that NetApp customers can take advantage of new SSD technologies as they emerge and add them to the FAS cluster, enabling longer storage system lifecycles.

## All Flash FAS as Primary Enterprise Storage

Another critical aspect to the addition of all-flash array availability to the FAS systems portfolio is the inclusion of clustered Data ONTAP feature/functionality for mission-critical enterprise data center environments. This includes data protection and business continuance functionality (SnapMirror, FlexClone, etc.), thin provisioning and OnCommand automation tools plus NFS and multiprotocol connectivity.

---

<sup>1</sup> See Evaluator Group Technical Insight paper entitled "Defining the Value of Modular Scale-out Storage."

To take maximum advantage of the inclusion of All Flash FAS, the latest release of Data ONTAP (version 8.3 and later) also demonstrates software-based optimization for flash including:

- Read path optimization for low latency resulting in a 270%+ increase in read I/O being serviced within a 1ms response time vs FAS system disk.
- Parallel threading of I/O that takes advantage of increased core densities in the All Flash FAS platforms and also results in low latency. I/O dispatched to flash storage is not single-threaded and serialized.
- Integrated data reduction technologies and SSD partitioning maximize flash utilization.
- Multiple options for data protection and business continuance are available including internal FAS system backup, backup to cloud, as well as full synchronous and asynchronous data replication. Secondary copies of data in flash can be maintained on disk storage within the cluster or in the cloud.
- Advanced Drive Partitioning provides a more efficient way of allocating usable data providing up to 17% additional usable flash capacity vs. previous ONTAP versions.
- In-memory write coalescing allows writes to be acknowledged to the hosts/client and coalesced in battery-backed NVRAM memory before being de-staged to flash. This prolongs flash module usable life.
- Pre-fetch and read-ahead algorithms pull data into FAS controller memory for faster access.
- WAFL blocks (random or sequential) are written into free areas of contiguous flash space, which allows any write to be sent to flash sequentially. Clustered Data ONTAP writes to free space whenever possible, minimizing overwrites for every dataset and maximizing flash media longevity.
- Because clustered Data ONTAP acknowledges writes after they are in DRAM and logged to NVRAM, solid state storage devices do not create delays in the write path. Therefore, write latencies are very low.

## Deploying All Flash FAS in High Visibility Application Environments

### Database and OLTP

Response time and transactional performance requirements are critical to production database environments. Flash as primary storage delivers response times measured in microseconds vs milliseconds for rotating disk yielding an order of magnitude increase in performance for database applications. And because the average relational database size is modest by unstructured data storage standards, multiple databases could be stored entirely in an All Flash FAS array.

High volume, transaction-oriented database applications can benefit from a number of All Flash FAS features mentioned earlier that maximize the cost efficiency of flash in database environments:

- For write I/O intensive workloads, writes are acknowledged after they are in DRAM and logged to NVRAM, before writing to flash.

- Data ONTAP always writes to a new location, eliminating overwrites.
- When high performance for an application is needed, the workload can be non-disruptively migrated from disk to flash within an ONTAP cluster<sup>2</sup>.
- Inline data compression and inline and always-on data deduplication reduces storage space required for databases and accelerates backup and recovery processes.
- NetApp Snapshot technology can be integrated with other database management interfaces (Microsoft SQL Server Virtual Device Interface for example) to create application-consistent snapshot copies with no downtime

Finally, the ability (mentioned earlier) to consolidate database workloads to fewer CPU cores allows database administrators to reduce software licensing costs that are charged on a per core basis. An example can be seen with Microsoft SQL Server 2014. NetApp reports increasing performance over a legacy disk array by approximately 4X while delivering space efficiency savings 17:1.<sup>3</sup>

### Virtualized Server and Virtual Desktop (VDI)

As a greater percentage of applications become hosted on virtualized servers, there are now more demands than ever on storage architectures to handle the performance of intensified, random I/O environments such as server virtualization and VDI. These increasingly concentrated workloads produced by tier 1 applications moving to virtualized environments are prime candidates for All Flash FAS.

In our experience, over 80% of performance issues reported by storage administrators dealing with virtualized environments are related to storage performance. As mentioned, the ability to handle today's tier 1 application environments within service level agreements (SLAs) can directly influence revenue generation on a per second, per hour, and per day basis. Therefore, there is a significant economic value in accelerating applications hosted on virtualized servers to increase the volume business operations, transactions and revenue. Additional capabilities of advantage to virtual server environments include:

- ONTAP data cloning technology supports very fast creation of writable copies of virtual machines.
- Using Storage Quality of Service (QoS), administrators can prioritize workloads in a multi-tenant, shared storage environment and avoid the over-consumption of storage resources by an I/O intensive application (a condition often called the "noisy neighbor problem").

---

<sup>2</sup> This is complementary to Oracle ASM which provides further load balancing of I/O across all LUNs or files in an Oracle ASM device group by distributing the contents of each data file evenly across the entire pool of storage in the device group.

<sup>3</sup> Test report available from NetApp in a report entitled "NetApp AFF8080 EX Performance and Server Consolidation with Microsoft SQL Server 2014," published May, 2015. Test results have not been validated by Evaluator Group.

- VMware vSphere VVOLs are supported to enhance the granularity of storage management in a VMware vSphere environment and implement automated data placement under vSphere.

Adding Virtual Desktop Infrastructure (VDI) workloads to the mix can result in an extremely high I/O environment that demands greater performance and lower response time vs rotating disk. Again, we note that the leading performance issue for VDI as reported by IT administrators is storage performance.

Virtual Desktop (VDI) administrators require significant levels of storage performance not only to counter the disruption of boot storms but to provide predictable and consistent response times in environments that include I/O-intensive desktops. All Flash FAS can meet these performance requirements and assure quality of service as experienced by virtual desktop users as the environment grows. In addition, cloning master OS images (gold images) for all desktops to read from in All Flash FAS provides high performance access to these gold images.

VDI environments are also price sensitive. Generally, the proposal of a project that converts physical to virtual desktops encounters resistance from upper level management when the cost of a virtual desktop exceeds that of a physical one. To respond, using deduplication in conjunction with All Flash FAS can provide significant space savings—and therefore cost savings—for data that needs flash performance because data in VDI environments is highly redundant. NetApp reports achieving a 10:1 reduction ratio and being able to provision storage at a cost of \$39 per virtual desktop.<sup>4</sup>

For deployment flexibility, FAS systems support VMware Horizon View; Citrix XenDesktop; and major hypervisors such as VMware vSphere, Microsoft Hyper-V, and Citrix XenServer, and Oracle VM and KVM.

### **Evaluator Group Assessment:**

*For customers looking to implement flash for primary data center storage, NetApp's All Flash FAS systems offer some significant advantages over competitive offerings. Customers can deploy All Flash FAS arrays for specific applications now and later extend them to private and hybrid cloud-based applications as enterprise IT's longer term strategy dictates. One investment serves both objectives.*

*The same can be said for existing Clustered Data ONTAP customers. They can simply add all-flash arrays to an existing storage cluster to accelerate the performance of database applications for example. They can do so non-disruptively while maintaining and adding to the Data ONTAP feature/function environment that supports existing applications and then extend ONTAP data fabric to cloud environments. Doing so also avoids the creation of yet another data silo to the enterprise storage environment as would be the case with an all flash array that has different operational and management characteristics.*

---

<sup>4</sup> Test results will be available in a NetApp Technical report available in July, 2014. Test results have not been validated by Evaluator Group.

*Enterprise-grade flash storage is not only an enabler of business-related gains. Flash is now becoming a critical part of the storage hierarchy and should be seen and implemented as a strategic data center investment. In addition, we note that Cloud Services Providers (CSPs) are challenging enterprise IT administrators to reduce the overall cost of IT. One way that CSPs compete with enterprise IT is on the basis of infrastructure efficiency. Flash allows enterprise IT to compete with CSPs on a more equal cost basis.*

## About Evaluator Group

Evaluator Group Inc. is dedicated to helping **IT professionals** and vendors create and implement strategies that make the most of the value of their storage and digital information. Evaluator Group services deliver **in-depth, unbiased analysis** on storage architectures, infrastructures and management for IT professionals. Since 1997 Evaluator Group has provided services for thousands of end users and vendor professionals through product and market evaluations, competitive analysis and **education**. [www.evaluatorgroup.com](http://www.evaluatorgroup.com) Follow us on Twitter @evaluator\_group

**Copyright 2015 Evaluator Group, Inc. All rights reserved.**

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, inconsequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.