

# All-Flash FAS: A Deep Dive



**Saurabh Modh**  
Product Manager



**Chetan Khetani**  
Senior Product Manager

The variety of all-flash storage arrays on the market—and the amount of hype surrounding some of the offerings—has reached bewildering proportions. In reality, there are really just four things by which all-flash arrays are measured: performance, storage efficiency, availability, and data management.

If you haven't heard much about our all-flash FAS storage systems yet, they stack up well to the competition in terms of performance and storage efficiency. When you factor in the proven availability of FAS hardware plus the ecosystem of data management features supported by our Data ONTAP® operating environment—it's a combination that's hard to beat.

In this article we "take the covers off" all-flash FAS. Specifically, we'll discuss what the all-flash FAS is (and isn't), how Data ONTAP is optimized for flash, how we deliver comprehensive storage efficiency, and how all-flash FAS performs for two popular all-flash use cases: VDI and database. We'll also talk a little about when to choose all-flash FAS versus our EF-series all-flash offerings.

## Introducing the All-Flash FAS

---

NetApp has been shipping all-flash FAS configurations for a number of years. However, with the release of the [FAS8000](#) early in 2014 and the recent release of the [FAS8080 EX](#), all-flash FAS has really come into its own. Built on the latest Sandy Bridge and Ivy Bridge CPUs, these platforms deliver the raw compute performance, I/O bandwidth, and high-speed network connectivity that all-flash arrays need. (If you're not familiar with the FAS8000 models, click the hyperlinks in this paragraph to learn more.)

An all-flash FAS system is just what it sounds like—FAS controllers configured with SSD shelves. You can have a FAS system with both SSD and HDD, but that isn't an all-flash FAS. Any FAS model can be configured as an all-flash FAS system, but the FAS8000 series is the sweet spot.

The great thing about the all-flash FAS is that it delivers all the capabilities that people have grown to expect from FAS hardware and Data ONTAP software. If you have an existing FAS cluster, you can drop in an all-flash FAS, non-disruptively migrate existing workloads, and immediately realize the benefits of all-flash. If you're an existing NetApp customer or are thinking about becoming one, all-flash FAS delivers all-flash performance with all the benefits of clustered Data ONTAP including:

---

### Explore

#### NetApp Voted 2014 Innovation Leader for All-Flash NAS

NetApp was selected by IT professionals as the leader in the 2014 All-Flash NAS category. Respondents were asked who they perceived to be the market, price, performance, reliability, service and support, and innovation leader. The result was announced along with other awards at the 2014 Flash Memory Summit.

[More](#)

---

- **Unified Storage.** All-flash FAS configurations are equally adept at serving both SAN and NAS workloads. Almost all other all-flash arrays are SAN only. NetApp was recently recognized by IT Brand Pulse as the leader in the 2014 all-flash NAS category.
- **Scale-Out with Data Mobility.** All-flash FAS configurations scale out up to 24 nodes, multiple petabytes of SSD capacity, and millions of IOPS. For some other leading all-flash arrays, scale-out exists mostly on the roadmap of future enhancements. You can scale out using just all-flash FAS, or you can mix all-flash and hybrid FAS systems and move data sets to all-flash without disruption when maximum performance is needed.
- **Nondisruptive Operations.** You can perform all maintenance, software, and hardware upgrades—including full tech refresh—without disrupting running applications.
- **Proven Availability.** All-flash FAS is based on the combination of proven FAS hardware (with hundreds of thousands of installed units in the field) and the Data ONTAP operating environment, which has been hardened against failures through 20 years of continuous development. Availability as measured in the field routinely hits six nines (99.9999%).
- **Quality of Service.** Built-in QoS lets you set performance limits on files, volumes, LUNs, or an entire storage virtual machine (SVM) to give priority to your most important workloads.
- **Comprehensive Storage Efficiency.** A full portfolio of storage efficiency technologies—including deduplication, compression, thin provisioning, cloning and others—increases effective capacity so you store more data with less storage. (More on this in a later section.)
- **Integrated Data Protection.** NetApp has spent years developing its data protection capabilities. All technologies are easy to manage, effective, space efficient, and deliver significant advantages for all-flash FAS. For example, you can save money by replicating expensive flash media to less expensive HDDs for disaster recovery.
- **Superior Application Integration.** Your storage needs to work with the applications and software tools you depend on. Data ONTAP is the number-one branded storage operating system, with broader and deeper application integration and a network of partners.

## Data ONTAP Is Optimized for Flash

---

You might hear claims that, because Data ONTAP existed before flash, it isn't designed to address the unique requirements that flash SSDs create. But the fact is that the Data ONTAP's Write Anywhere File Layout (WAFL), scale-out capabilities, and feature set are ideal for all-flash environments. All-flash FAS delivers predictable low latency. Not only is average latency low (1 millisecond), but the variability in latency is minimized. This high level of consistency benefits applications.

We provide tools for comprehensive [performance monitoring](#) to help you get the most from your investment, and Data ONTAP contains many optimizations that make it particularly well suited for flash.

### Lighting Up All the Cores

There's no question that it takes a lot of computing horsepower to drive all-flash workloads. As we've already seen, the FAS8000 series delivers this raw horsepower. We've done a tremendous amount of work in Data ONTAP to make sure we fully utilize all available cores in multi-core processors with unrestricted scheduling of processes across all cores. This improves latency in the CPU complex and allows us to simultaneously support both front-end workloads and back-end data management tasks such as backup and replication.

CPU0	CPU1	CPU2	CPU3	CPU4	CPU5	CPU6	CPU7	CPU8	CPU9	CPU10	CPU11	CPU12	CPU13	CPU14	CPU15	CPU16	CPU17	CPU18	CPU19
97%	98%	97%	97%	97%	97%	98%	97%	97%	97%	97%	97%	98%	97%	97%	98%	97%	97%	97%	98%
96%	97%	97%	96%	96%	97%	97%	96%	96%	96%	96%	97%	96%	96%	96%	97%	97%	97%	97%	96%
97%	97%	98%	97%	97%	98%	97%	97%	98%	97%	97%	98%	97%	97%	97%	98%	98%	97%	98%	97%
97%	97%	97%	97%	97%	97%	97%	97%	97%	97%	97%	97%	97%	97%	97%	97%	97%	97%	97%	97%
97%	97%	97%	98%	97%	97%	98%	97%	97%	97%	97%	98%	97%	97%	97%	97%	98%	98%	98%	98%
96%	97%	96%	96%	96%	96%	96%	97%	96%	96%	97%	97%	97%	97%	97%	97%	97%	97%	97%	97%
97%	98%	98%	98%	97%	98%	98%	97%	98%	97%	97%	98%	98%	98%	98%	98%	98%	98%	97%	98%
96%	97%	97%	96%	97%	97%	97%	97%	97%	97%	96%	97%	96%	97%	97%	96%	97%	97%	97%	97%
97%	97%	97%	97%	97%	97%	97%	97%	97%	97%	97%	97%	97%	97%	97%	97%	98%	97%	97%	97%
97%	97%	96%	97%	97%	97%	97%	97%	96%	97%	97%	97%	97%	97%	97%	97%	97%	97%	97%	97%

**Figure 1)** CPU utilization for an all-flash FAS8080 EX running clustered Data ONTAP 8.2.2. Cores are evenly utilized at near maximum load during database performance testing.

Using cores efficiently translates to more IOPS per core. This is born out by the fact that the [FAS architecture delivers more than 5X the IOPS/core of a competing scale-out platform](#). (Both platforms used hybrid storage configurations for these benchmarks.)

### Optimized Write Performance

Random write operations affect the usable life of SSDs. It's a well-known fact that the erase-before-write behavior of flash can produce undesirable effects such as [write amplification](#), which can reduce write performance and shorten device life. Data ONTAP contains a number of optimizations that enhance write performance—for both SSDs and HDDs.

- **NVRAM Accelerates Writes.** Incoming writes are processed in memory, protected in NVRAM, and acknowledged immediately to the client or host, resulting in extremely fast write response time. Writes to SSDs are not on the critical latency path.
- **Optimized writes.** All writes are coalesced in memory before being destaged to SSDs. Write metadata is written alongside user data to minimize the number of operations required to commit data to SSDs. Write striping across SSDs results in maximum performance and reduces uneven SSD wear.
- **No Block Overwrites.** When an overwrite is received, the block is always written to a new location. Existing blocks (data and metadata) are never re-written to the same location.

This combination of features reduces the impact of random writes to SSDs in all-flash FAS:

- SSDs have to move less data when they do garbage collection
- The "effort" to achieve a given level of write-throughput is reduced
- Write amplification is lower and SSD life is longer
- System-level performance is improved
- Reduced write effort results in more consistent read latencies for mixed read/write workloads

Because of these optimizations, NetApp is able to offer an SSD warranty that is not tied to any wear-leveling limits, and as you'll see later, all-flash FAS performs very well on write-intensive workloads such as VDI.

### Exceptional Data Integrity Features

Another characteristic of flash is that it may be subject to bit disturbances and losses over time as flash cells age. Data ONTAP provides comprehensive protection to prevent data loss and ensure integrity:

- Dual-parity RAID protection

- In-memory and on-disk block-level checksums
- Comprehensive lost write detection and correction to avoid silent data corruption
- SSD-level bad block management
- Continuous media scrubbing to detect and correct errors

## Comprehensive Storage Efficiency

---

If we had to pick the one aspect of the all-flash market where the hype has reached a fever pitch, it would be storage efficiency. Because flash storage capacity is expensive relative to HDDs this makes sense: The more capacity you can "save" through storage efficiency, the less money you spend on storage.

At the risk of sounding as if we're making inflated claims of our own, it wouldn't be exaggerating to say that NetApp pioneered many of the ideas behind modern storage efficiency. From the beginning, we've been developing technologies that let you store and protect more data; it's certainly part of what has made Data ONTAP the number-one branded storage operating environment.

Unless this is the first Tech OnTap® article you've ever looked at, you're probably already well aware of NetApp® storage efficiency. A fairly complete list is shown below (most have links to past Tech OnTap articles that describe them):

- Technologies that boost utilization and reduce primary storage requirements
  - [Thin provisioning](#)
  - [Deduplication](#)
  - [Compression](#)
  - [Cloning](#)
  - [RAID-DP](#)
- Technologies for efficient data protection
  - Snapshot™ copies
  - [Mirroring](#)
  - [Disk-based backup](#)
  - [SnapManager® suite](#)
  - [SnapProtect®](#)

(Also check out [Storage Efficiency in Clustered Data ONTAP](#) for the latest on all of these technologies.) The key thing is that all these technologies are designed to work together so the effects are additive. For example, if you replicate a volume that's deduplicated and compressed, the replica is in the same state.

While we hold ourselves to a higher standard than some of our competitors—who may quote numbers based on a handful of selected customers—we're confident that all-flash FAS can deliver an effective capacity of 5:1 on average. This is based on data from hundreds of thousands of installations over many years. Of course, as is always the case, the results you achieve will depend on the characteristics of your data sets and workloads. Virtualized environments typically exceed this number (10:1 is common) while database environments come in lower (2:1).

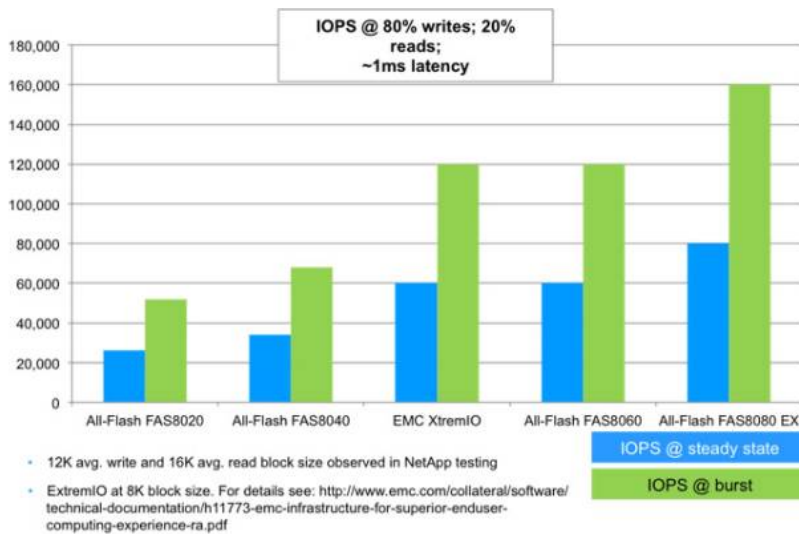
## All-Flash FAS Use Cases

---

Two major use cases for all-flash FAS are virtual desktop infrastructure (VDI) and database. These two also illustrate the capabilities of all-flash FAS, because VDI is write intensive, while database workloads are often tilted toward read performance.

### VDI

We've done extensive VDI testing of the entire FAS8000 product line, measuring both steady-state and burst performance based on a workload that's 80% writes and 20% reads. As you can see in Figure 3, there's a nice distribution between the four models. The graph also shows a well-known competitor, XtremIO, for reference. As you can see, FAS8060 performance matches that of XtremIO and the FAS8080 EX exceeds it by a 30% margin.



**Figure 2)** All-flash FAS VDI performance versus a competitor.

Our testing used an average write size of 12KB and an average read size of 16KB versus 8KB block sizes used by competitors. The different block size is a result of our testing methodology. We used Login Virtual Session Indexer (Login VSI), the industry-standard load-testing tool for testing the performance and scalability of virtual desktop infrastructure. You can read complete details of our VDI testing for the FAS8060 in [TR-4307: NetApp All-Flash FAS Solution for VMware Horizon View](#).

Based on that testing, the FAS8060 has enough performance headroom to support 4,000 desktops (with HA failover) and [comes in at \\$55/desktop](#). That's an extremely competitive price for VDI.



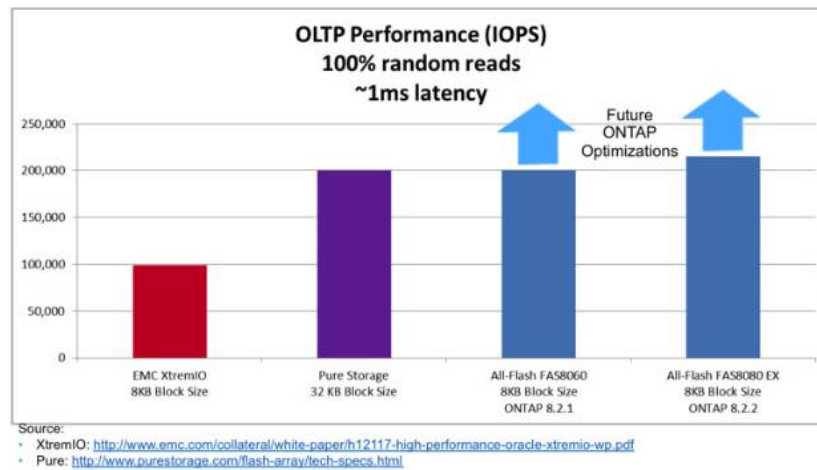
While benchmarks are valuable, it's also important to understand why real VDI customers choose all-flash FAS:

- A large insurance and investment firm picked all-flash FAS for its Citrix VDI environment after an evaluation of major all-flash players. In the end they chose all-flash FAS because of its proven performance combined with the integrated data management capabilities of clustered Data ONTAP including SnapMirror® for replication, FlexClone®, and the ability to do nondisruptive DR testing.
- A clinical service provider chose all-flash FAS for all its applications including its VDI environment and electronic medical records (EMR). For them, faster access to data translates directly to better patient care.

### Database

For database, we did performance testing using a workload that is 100% random reads. Performance is shown in Figure 4, and includes two well-known competitors. The FAS8060 matches the best performance of the two, while the FAS8080 EX exceeds the performance of both. Upcoming enhancements in Data ONTAP are expected to

further boost read performance by as much as 50% depending on workload.



**Figure 3)** All-385 FAS database performance relative to public numbers from several competitors.

The table below provides a little more detail. As you can see, the FAS8080 EX compares favorably with competitors in terms of \$/IOPS and \$/GB. By these measures it is highly competitive, without even factoring in the proven availability of the platform or extensive data management features.

**Table 1)** Comparison of all-flash FAS8060 with competitors for database workloads.

HW Configs	All-Flash FAS8060	XtremIO	Pure Storage
Usable Capacity	8.7 TB	7.5 TB	6 TB
IOPS (100% read)	200,000 <sup>1</sup>	100,000 <sup>2</sup>	200,000 <sup>3</sup>
Net \$/IOPS (100% read)	\$0.97	\$1.94	0.96
Net \$/GB (Usable)	\$22	\$26	\$32
Net \$/GB (Effective <sup>4</sup> )	\$11	\$13	\$16

<sup>1</sup>FAS8060 with ONTAP 8.2.1 results. Block size is 8KB  
<sup>2</sup>Block size is 8KB  
<sup>3</sup>Block size is 32KB (Source: Competing Startup Spec sheet)  
<sup>4</sup>Storage Efficiency ratio of 2:1 for DB workloads

Of course, database customers find their own reasons for choosing all-flash FAS:

- A major hotel chain chose all-flash FAS for its reservation system to meet the performance requirements of its new online reservation system. all-flash FAS was selected for its backup and data management advantages.
- A government customer needed to speed up SQL Server. Every day they ingest 6TB of data into a data warehouse, so throughput is as important as IOPS for them. all-flash FAS testing showed it could deliver 4.4GB/s with two nodes, nearly 3 times the measured performance of a top competitor. They also needed a DR copy and several development copies, which could be provided easily using SnapMirror and FlexClone with minimum storage (and network) overhead.
- An international bank picked all-flash FAS to support both VMware® and high-performance SQL Server applications. As an existing NetApp customer, they faced a simple choice because they already recognized the value of SnapVault® and SnapMirror. Overall, the bank saw an 8X improvement in overall application responsiveness as a result of the change.

## Making the Right All-Flash Decision

At this point, it should be pretty clear that the all-flash FAS lineup does a great job performing all the functions of an all-flash array, delivering a combination of great performance, comprehensive storage efficiency, proven availability, and extensive data management. When combined in a cluster with other FAS configurations, you can

easily migrate workloads to and from all-flash as your needs change, or provide secondary data copies on less expensive HDD storage.

But the all-flash FAS is not the only arrow in our quiver. NetApp also offers the EF-Series all-flash arrays. The [EF540](#) and [EF550](#) have been described previously.

The EF-Series is a great choice for I/O-intensive workloads where data management is being handled by the application. For instance, in many Oracle environments backup and DR are handled by Oracle itself and the all-flash array simply needs great performance and reliability at a great price. The EF550 sustains in excess of 400,000 IOPS at submillisecond latencies and up to 12GB/sec throughput.

All-flash FAS performs extremely well in VDI and database environments and is a great choice in situations that can benefit from storage efficiency, NAS, and the data management capabilities of Data ONTAP. Existing FAS customers will find that it is managed and operated the same as any other FAS system, making adoption painless.



**By Saurabh Modh, Product Manager and Chetan Khetani, Senior Product Manager**

Saurabh is a product manager in the NetApp FAS Systems Group responsible for all-flash FAS, FAS3200 and FAS6200 storage systems. Prior to joining NetApp, he worked as a Project Lead at Oracle doing software development for Oracle Enterprise Manager. Saurabh received his MBA from the Foster School of Business at the University of Washington.

Chetan is responsible for all things flash in Data ONTAP. This includes both all-flash FAS and hybrid solutions, which have seen a 10X increase in recent years. He is a frequent contributor in the NetApp Executive Briefing Center, where he was recently chosen Speaker of the Quarter for his efforts. Before joining NetApp, Chetan spent 11 years at Cisco where he worked on technologies for network virtualization, deep packet inspection, and context-based infrastructure. He holds an MBA and an MS in management information systems.

---

### *Quick Links*

- › [Tech OnTap Community](#)
  - › [Archive](#)
  - › [User Groups](#)
  - › [PDF](#)
-