

# Storage Design: This is Not Your Father's Storage Subsystem



**Tushar Routh**

Storage Hardware Product Manager

It seems that every few years a change comes along that causes a fundamental shift in the way we think about and deploy storage media—and, more often than not, NetApp is leading the way.

Back in 2002, NetApp pioneered the use of capacity-oriented drives for secondary storage instead of tape. A few years later—in large part due to the increased reliability made possible by our RAID-DP® technology—NetApp made it feasible to deploy this type of drive for primary storage for the first time.

In the last three years, the disk drive industry has undergone a significant shift away from the Fibre Channel interface to SAS. NetApp was one of the first major storage vendors to adopt SAS and a clear leader in designing storage solutions with SAS technologies. Today, widespread acceptance and adoption of flash technology are creating yet another change.

The result of all these developments is that you have a lot more options and, although having options is clearly a good thing, it also means that there's more to think about when you're trying to choose the "best" storage media or optimize a storage subsystem for a particular application workload.

In this article I explore some near-term storage trends as well as take a look at what might happen a little further down the line. I also provide some guidance regarding what NetApp believes are the storage subsystem best practices that will keep you ahead.

## Near-Term Media Trends

---

With the availability of flash, deploying storage is no longer just about getting enough spindles to support your workload; increasingly it's about pairing the right disk drives with the right flash options. Combining hard disk drives (HDDs) and flash to create hybrid storage really changes the calculus of storage deployment. This is one of the key reasons why NetApp offers an expanded range of drive options. You can mix various types of media in a single FAS storage system or a single cluster to address a wide variety of storage requirements.

---

## Explore

### Recent NetApp Storage News

Recent blogs that focus on new NetApp storage options include:

- [New Drives Enhance Density and Flexibility](#)
- [New 1.2TB HDD Expands Options for Transactional Workloads and VST Deployments](#)
- [NetApp First with Optical SAS for Flexibility and Resilience](#)

Storage design blog series:

- [HDD](#)
- [SSD](#)
- [Shelves and Interconnects](#)
- [Guidelines](#)
- [NetApp Storage Encryption](#)

### Recent Blogs on Security and NSE

Security is a hot topic. NetApp, in conjunction with our security partner SafeNet, is working to keep you up to date.

- [For Your Eyes Only: How to Use Encryption to Keep Big Brother—and Thieving 3rd Cousins—Away from Your Data](#)
  - [Keys to Security Success](#)
  - [Storage Design Series: NetApp Storage Encryption](#)
-

Let's start by taking a look at what's happening today and what will happen in the near term with four classes of drives:

- High-capacity HDDs
- Performance HDDs
- Self-encrypting HDDs
- Ultraperformance SSDs

## High-Capacity HDDs

In order to deliver maximum density and the lowest cost per gigabyte of capacity, high-capacity HDDs continue to use the large-form-factor (LFF) 3.5" format, typically with a rotation speed of 7,200 RPM.



Near-line SAS disk drives—which combine the SAS interface with the media and rotation speed of enterprise SATA drives—are becoming the preferred choice for near-line and production workloads. Over the next year or two, we should see available options reach 5TB and 6TB capacity points. The SATA interface will remain the preferred choice for backup and archival workloads over that period.

**NetApp High-Capacity Options.** Late last year, NetApp released a [4TB high-capacity drive](#) for use in our 4U, 48-drive DS4486 high-capacity disk shelf, making NetApp the first major storage vendor to ship HDDs at this capacity. The combination of 4TB drives and the DS4486 disk shelf provides a level of capacity and density that makes it ideal for near-line, archive, and backup applications. In April 2013, NetApp also released a separate 4TB drive for use in our [4U, 24-drive DS4246 disk shelf](#) to address high-capacity storage needs for production workloads. Our full portfolio of high-capacity drives currently includes 1TB, 2TB, 3TB, and 4TB options.

## Performance HDDs

The performance HDD market has moved away from 3.5" LFF HDDs to 2.5" small-form-factor (SFF) options. The major HDD suppliers have signaled that 3.5" 15K RPM LFF drives will go out of production in the near future. NetApp introduced its 2U, 24-disk DS2246 disk shelf some time ago in preparation for this transition.

In the past, NetApp was primarily concerned with delivering low capacity points for this class of drive to allow you to deploy the number of spindles needed to meet performance requirements while minimizing overprovisioning of capacity. However, now that it's become common to use a combination of performance disks and flash—either NetApp® Flash Cache™ or Flash Pool™ intelligent caching—to address performance, we've begun making larger capacities available.

With more capacity points, you can deploy Flash Cache or Flash Pool with an optimum number of spindles to address the capacity and performance needs of your workloads. For example, suppose a particular application needs 20TB of capacity and 10K IOPS and has a cache hit rate of 80% with Flash Cache. The 1.2TB drives will deliver more than enough spindles to satisfy cache misses with good performance.

For existing workloads running on NetApp, you can use [predictive cache statistics](#) to determine what your cache hit rate will be for a given amount of flash. This helps dial in the right amount of flash and HDDs before you invest in new media.

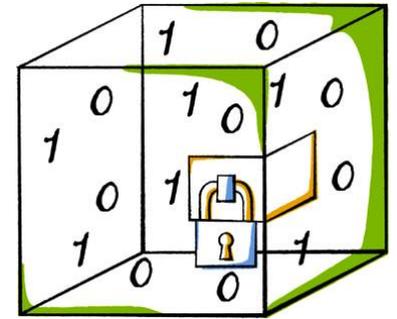
**NetApp Performance Drive Options.** We [recently released a new 1.2TB 10K RPM HDD](#) that expands the range of SFF hard disk drives we offer. The full portfolio now includes 450GB, 600GB, 900GB, and 1.2TB capacities. Even larger capacities are likely to come in the future with drive manufacturers due to deliver 1.8TB SFF capacities in the next year or so. Over time, some of the lower capacity points will be phased out. There are 15K SFF drive options on the market as well; however, NetApp does not offer them because they cost two to three

times more than 10K SFF drives. Because SSDs offer better I/O density, we believe that they are a better option than 15K SFF HDDs and are likely to displace them.

## Self-Encrypting Drives

With all the recent news about corporate and government espionage, there's been a definite uptick in interest in security and encryption solutions. NetApp Storage Encryption (NSE) is the NetApp implementation of full-disk encryption (FDE) using self-encrypting drives from leading drive vendors.

Because encryption and decryption take place on the drive itself after data is written by Data ONTAP or as part of the disk read, NSE operates seamlessly with features such as deduplication and compression. All data on a drive is automatically encrypted, so using NSE is an easy way to protect data at rest while maximizing the ROI of your NetApp storage.



As you might expect, this technology is front and center for government, healthcare, and financial organizations. The physical drives themselves are tamperproof, and NSE prevents unauthorized access to encrypted data at rest. It prevents someone from removing a drive or shelf of drives and mounting and accessing them elsewhere. In addition, it prevents unauthorized access when drives are returned after a drive failure and simplifies the disposal of drives.

All FDE drives that NetApp sells adhere to the Trusted Computing Group AES-256 encryption standard. We also require FIPS 140-2 certification, which is a standard requirement for the public sector.

NSE adds some expense to a storage deployment. You can't mix encrypting and nonencrypting drives in the same platform, and the drives have additional cost due to added components and the requirement that the physical drives be tamperproof. In addition, key management—which is provided by an external device—adds expense.

Key management for NSE is provided by an external solution. NetApp has partnered with SafeNet to offer the [SafeNet KeySecure Key Manager](#), available directly from NetApp and from our partners.

**NSE Options.** Our NSE drive portfolio includes a 600GB and a 900GB performance drive and 3TB and 4TB LFF high-capacity options. We plan to offer a self-encrypting SSD in the coming months, so ultimately—although a release date has not been set—we plan to support Flash Pools that are fully encrypted after we have a self-encrypting SSD in our portfolio.

## Solid-State Drives

Capacities of SSDs have been growing fast in recent years. Although this growth will continue, the flash memory used in SSDs is running up against the same lithography limits as other types of semiconductors. The path to increased capacity for flash has been to shrink the feature size on each chip to deliver more capacity per chip. Current NAND flash devices are using a 2Xnm-class process (20-29nm feature size) and rapidly moving to a 1Xnm process (10-19nm feature size). At the same time, enterprise SSDs have made the transition from single-level cell flash components to multi-level cell devices.

Today, you can use SSDs either as persistent storage—like any other type of drive—or as part of a Flash Pool that combines HDDs with SSDs to accelerate random reads and writes. Here are some things to keep in mind for each type of deployment.

- **For persistent storage use RAID-DP.** When creating SSD aggregates for persistent storage, the best practice is to use RAID-DP for SSD RAID groups.
- **For Flash Pools use RAID 4.** As of Data ONTAP 8.2 you can mix RAID types in a Flash Pool. This allows you to use RAID-DP for the HDDs within a Flash Pool while using RAID 4 for the SSDs, reducing the cost for a given amount of usable flash.



NetApp likes to say that Flash Pool combines the capacity of HDD with the performance of flash, but it's really more than that. SSDs provide the most benefit for random, transactional workloads. HDDs are actually quite good at sequential workloads, especially on a cost basis. Combining the two types of media in a single aggregate lets you benefit from SSDs for their transactional performance and HDDs for sequential performance without having to know the exact I/O behavior of every workload when you architect the storage.

To learn more about Flash Pool and all of NetApp's flash options, including the [EF540 flash array](#), check out the recently released book [Flash for Dummies](#).

**NetApp SSD Options.** As with HDDs, NetApp is introducing new SSDs on an aggressive schedule. We currently offer 200GB and 800GB capacities and will add more options in the coming months to provide a broader range of choices for even greater storage optimization.

## Maintenance Center

To help you get the greatest reliability from each HDD, Data ONTAP includes Maintenance Center, which performs proactive health monitoring of drives, distinguishes between transient events and real underlying issues based on drive diagnostics, and attempts preventive maintenance when necessary.

Maintenance Center automatically manages disk failures through a systematic failure-verification process of the failing disk without removing it from your storage system. When a disk is identified as a potential failure, Maintenance Center takes over. Data is migrated from the disk onto a spare through Rapid RAID recovery (it copies data directly from the disk before failure occurs) or reconstruction. The process occurs without user intervention. If transient errors can be repaired, the disk is returned to the spares pool.

The functions of Maintenance Center do not apply to SSDs since SSDs don't experience the same types of transient failures as HDDs. However, Data ONTAP tracks the usage of SSDs and flags any drives that are reaching the end of their usable lives.

## Disk Shelves

---

NetApp currently offers a line of four disk shelves for FAS systems. All of these shelves are accessible from the front for easy serviceability, regardless of their location in the rack, and all are designed to be highly reliable with no single points of failure. On all shelves, shelf firmware upgrades are nondisruptive and alternate control path provides out-of-band management.

## Disk Shelf Options

**DS2246.** The DS2246 is our performance-optimized shelf that packs 24 drives in only 2U of rack space using SFF drives. Compared to the DS4243 disk shelf, the DS2246 doubles the storage density, increases performance

density (IOPS per rack unit) by 60%, and reduces power consumption by 30% to 50%.

**DS4246.** The DS4246 provides an ideal balance between performance and capacity. It is 4U high and supports 6Gb/sec SAS connections. It can be configured with either 24 LFF high-capacity disk drives or a combination of SSDs and high-capacity disk drives to support Flash Pool configurations.

**DS4243.** The NetApp DS4243 is 4U high and supports up to 24 hard disk drives (high capacity or high performance) with a 3Gb/sec SAS connection.

**DS4486.** The capacity-optimized DS4486 holds 48 high-capacity disk drives. This disk shelf looks like the DS4246 from the front, but it is slightly longer and uses a tandem disk carrier to enclose twice as many LFF disk drives in 4U of rack space. In contrast to many capacity-optimized disk shelves, the DS4486 can be serviced from the front, and 10 DS4486 shelves in a 42U rack weigh less than 2,000 pounds (910kg). The rack can be supported by a raised floor in a traditional data center.

The DS4243 will be discontinued in favor of the DS4246 in the near future as the LFF performance HDDs it supports are phased out. Otherwise, we're pretty happy with this lineup of disk shelves and don't anticipate any major changes in the next few years.

**Table 1)** Comparison of NetApp disk shelves for FAS/V-Series storage systems.

Specification	 DS2246	 DS4246	 DS4243	 DS4486
Rack units	2U	4U	4U	4U
Drives per shelf enclosure	24	24	24	48
High-capacity HDDs		✓	✓	✓
Performance HDDs	✓		✓	
Self-encrypting HDDs <sup>1</sup>	✓	✓	✓	
Ultrapformance SSDs	✓ (pure and mixed) <sup>2</sup>	✓ (mixed only) <sup>2</sup>		
I/O modules	Dual 6Gb/s	Dual 6Gb/s	Dual 3Gb/s	Dual 6Gb/s
Drive carrier form factor	2.5" Small form factor	3.5" Large form factor	3.5" Large form factor	3.5" Large form factor
Drive carrier	Single drive	Single drive	Single drive	Tandem drives

<sup>1</sup>Self-encrypting HDDs adhere to standards such as AES-128, AES-256, and FIPS 140-2.

<sup>2</sup>A "pure" SSD shelf contains SSDs only; a "mixed" shelf contains a combination of SSDs and HDDs for use by Flash Pool. Flash Pool also works in "shelf-to-shelf" configurations in which SSDs from a pure SSD shelf are combined into an aggregate with HDDs from other shelves.

<sup>3</sup>For details about MetroCluster, go to <http://www.netapp.com/us/media/ds-2893-metrocluster-solnbrief.pdf>.

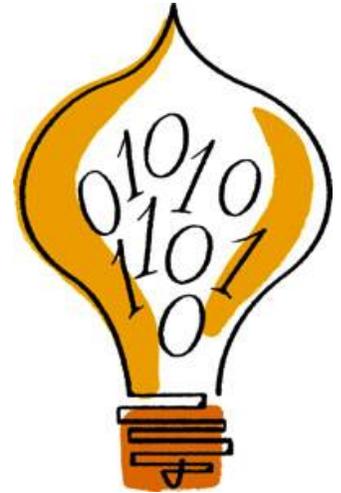
## Interconnects

---

NetApp was one of the first storage vendors to make the transition from Fibre Channel to SAS as a disk

interconnect, and we believe that SAS offers significant benefits in that role for the foreseeable future. The 6Gb/second SAS-2 connections we use today still deliver more than adequate bandwidth, and SAS-3 (12Gb/sec) is coming.

The only limitation of SAS is the relatively short run length of the standard copper cables. NetApp recently announced a [family of optical SAS products](#) to address this problem—the first company in the storage industry to do so. Optical cabling will address two main issues:



- **Cabling limitations.** In a crowded data center it can be tough to stay within the 20M SAS cable limit. With much longer cable runs, optical SAS fixes this problem; you can add new disk shelves wherever you have rack space available.
- **MetroCluster.** MetroCluster™ technology lets you synchronously mirror data at campus and metropolitan area distances for continuous data availability. NetApp optical SAS cabling lets you create a "zero-footprint" MetroCluster configuration that spans up to 500M using preexisting optical infrastructure without needing SAS to Fibre Channel bridges that add complexity and expense.

Be sure to check out the [SAS Storage Cabling and Infrastructure FAQ](#) for more information.

## Longer-Term Trends

---

In the longer term, NetApp is tracking a number of additional media trends. Certain classes of disk drives are approaching the cost per GB of tape drives, opening the door for backup and archival based exclusively on disk.

- **Hybrid drives** that combine HDD and flash technology in a single device are starting to appear in the market. The use case for these is currently limited to certain server environments.
- A new class of HDD for **archival** is in development, offering a lower price point and a very low duty cycle. These HDDs will be used to archive data for long-term retention and compliance, addressing traditional WORM—what I jokingly refer to as "Write Once, Read Maybe"—requirements. These drives will require spin down when not being accessed to save energy and extend the lifetime of the device. As archival drives become a reality, it may be unnecessary for the disk shelves housing them to have dual power supplies, redundant paths, and so on, so we may consider a new, lower-cost shelf design.
- Another new class of drive is targeted for the **cloud or big data**. These drives also have a lower duty cycle and are basically an offshoot of current near-line technology. The drives are used primarily in environments that maintain three copies of data for redundancy. For most applications in which this level of redundancy is not required we see fewer spindles plus RAID as a cheaper alternative. NetApp is looking at this class of drive for **backup storage** in combination with our dense DS4486 shelf.
- **Consumer-grade disk drives** are being talked about as a cheaper alternative to enterprise disks, but we don't currently see this becoming a reality because the failure rates remain so much higher.

On the solid-state front, although there's no question that SSDs are changing the game when it comes to storage subsystem design, flash memory has some limitations. Flash cells can endure only a limited number of write cycles before they wear out. Newer technologies such as phase-change memory and resistive RAM are being discussed as a way to overcome these limitations, but it's too early to tell which if any of these technologies will emerge. It's clear that solid-state storage in some form will continue to have an expanding role in the overall storage market, especially as the economics improve.

## Guidelines for Storage Subsystem Deployment

---

With all the recent attention to flash storage, it would be easy to conclude that HDDs are going to be eclipsed in short order. There are certainly use cases for dedicated all-flash arrays, which is why NetApp released the [EF540 flash array](#) last year and why we're busily working on the [FlashRay™](#) storage array, our next-generation scale-out all-flash offering.

However, HDDs—especially in combination with flash—aren't going away any time soon. The economics of HDD and HDD/flash solutions relative to all-flash solutions still make them the sweet spot for a lot of storage workloads.

Here's what NetApp suggests.

- For applications that need the most consistent performance with the lowest latency, choose the EF540 flash array or a FAS system with SSDs.
- For other high-end workloads, deploy FAS with performance disks in combination with one (or more) of our flash-based caching technologies: Flash Cache (storage controller level) or Flash Pool (disk subsystem). NetApp put a lot of time and effort into providing options to let you put flash where you need it. (You can find more guidance on choosing among flash options in our recently released [Flash Storage for Dummies guide](#), and this [white paper](#).)

The DS2246 disk shelf is the best choice for this type of deployment because it can accommodate performance HDDs, SSDs, or a combination (for Flash Pool deployments) and because it offers superior performance density.

- For near-line or capacity-oriented workloads, deploy capacity disks in combination with flash. The DS4246 disk shelf is a good choice here because it supports both high-capacity HDDs and combinations of HDDs and SSDs.
- For backup or archival workloads, deploy capacity disks. The DS4486 disk shelf delivers maximum capacity per rack unit.
- For purely sequential workloads, capacity or performance disks deliver good performance at a lower price than that of SSDs.
- If you're not sure what the I/O characteristics of your workload are, or you need to support a variety of workloads that may include both transactional and sequential I/O patterns, HDD plus flash options are once again a good bet.

Here are a few best practices to keep in mind.

- Except for Flash Pool, don't mix different media types in the same SAS stack.
- A single FAS system can support multiple aggregates with different media types (performance HDD, high-capacity HDD, SSD) to address the needs of different workloads, or you can deploy clustered Data ONTAP and dedicate specific cluster nodes for specific media and workload types.
- Use RAID-DP (with the possible exception of SSDs deployed for Flash Pool).
- Make sure that RAID scrubs are turned on for RAID groups containing HDDs (this is the default setting) to keep your drives healthy.
- When deploying storage for which length limitations are likely to be an issue or when deploying "stretch" MetroCluster, choose optical SAS for easier cabling.
- Follow the guidelines in the recently updated [Storage Subsystem Resiliency Guide](#) to achieve maximum resiliency.

A little upfront planning—and consideration of the guidelines I've outlined here and in previous posts—will go a long way in helping you choose storage that is best suited for your particular workloads and needs. As you've no doubt noticed, there are a lot of options to consider: HDD versus SSD, SSD deployed as cache or persistent storage, different capacity points, and so on. Of course, NetApp experts as well as our worldwide network of

partners are available to assist you in your decision-making.

---

› **Got opinions about  
Storage Design?**

Ask questions, exchange ideas,  
and share your thoughts online in  
NetApp Communities.



**By Tushar Routh, Storage Hardware Product Manager**

Tushar has been at NetApp for more than three years and is responsible for FAS storage and RAID, including HDD, SSD, drive enclosures, and related infrastructure. He has over 20 years of industry experience, including more than 10 years working in the HDD manufacturing divisions of IBM and Hitachi.

---

*Quick Links*

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