



White Paper

Deploying Flash in the Enterprise

Choices to Optimize Performance and Cost

Paul Feresten, Mohit Bhatnagar, Manish Agarwal, and Rip Wilson, NetApp
April 2013 | WP-7182

Executive Summary

Flash is quickly emerging as the preferred way to overcome the nagging performance limitations of hard disk drives. However, because flash comes at a significant price premium, outright replacement of HDDs with flash only makes sense in situations in which capacity requirements are relatively small and performance requirements are high. Deployment approaches—including hybrid storage arrays, server flash, and all-flash arrays—that combine the performance of flash with the capacity of HDDs can be cost effective for a broad range of performance requirements.

NetApp offers a full range of flash solutions, including server flash, hybrid storage arrays, and all-flash arrays. We've done a careful analysis of the cost of each solution at various combinations of performance and capacity to help you understand how to choose the best solutions to address your storage challenges based on your performance needs (IOPS and latency), capacity requirements, working set size (amount of hot data), budget, and data protection objectives.

TABLE OF CONTENTS

1	Introduction	3
2	Deploying Flash for Enterprise IT	4
3	A Complete Flash Portfolio for Enterprise Needs	5
4	Hybrid Storage Solutions	6
4.1	Data ONTAP and the Virtual Storage Tier	6
4.2	E-Series Solutions	7
5	Server Flash	8
6	All-Flash Arrays	9
6.1	The EF540	9
6.2	FlashRay	10
7	Conclusion	10

LIST OF FIGURES

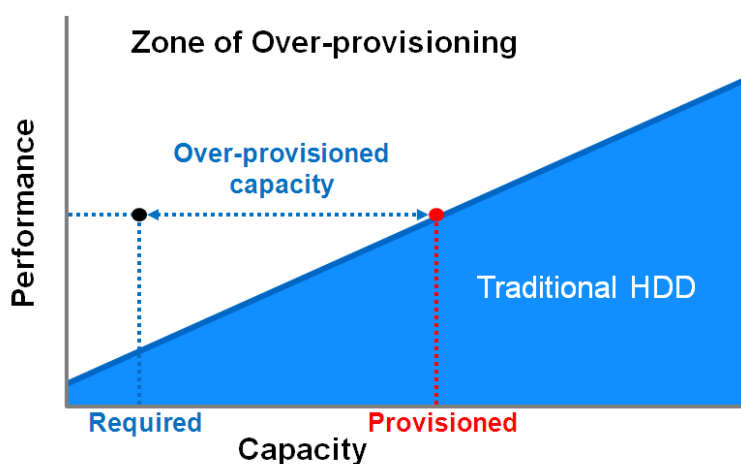
Figure 1)	HDD performance versus capacity.....	3
Figure 2)	Enterprise storage options	4
Figure 3)	Performance/capacity analysis.....	4
Figure 4)	The NetApp flash portfolio.....	5

1 Introduction

Flash technology is changing the way that enterprises approach storage. After years of use in the consumer market, flash has reached a price point and level of maturity at which it is being actively deployed to address the needs of business-critical applications.

Hard disk drives (HDDs) have some nagging deficiencies that make provisioning storage for applications with high-performance demands difficult. Because HDDs are capable of performing no more than 300–400 random I/O operations per second (IOPS), a storage system capable of delivering tens of thousands of IOPS requires hundreds of disks—even when the capacity is not needed. Overprovisioning disks to achieve performance goals is a significant capital expense and wastes rack space, power, and cooling. High-performance workloads increasingly require 100,000 IOPS or more, further exacerbating the problem.

Figure 1) HDD performance versus capacity. Overprovisioning is necessary whenever the required performance cannot be achieved at the required capacity.



The fastest HDDs have access times of 3–4 milliseconds, resulting in latencies much slower than flash-based SSDs, which have latencies measured in microseconds and perform thousands of IOPS per device. HDDs alone may no longer meet the needs of latency-sensitive applications.

Because of clear performance advantages coupled with significantly lower power consumption, flash SSDs and other flash devices are beginning to take the place of high-performance HDDs. However, because SSDs currently cost more than 10 times as much per GB of usable storage, IT teams are still searching for the best strategies to deploy flash technology to deliver performance where it's needed while minimizing overall storage costs.

NetApp has been shipping flash to accelerate enterprise applications since 2009, well before many current flash companies came into existence. As of April, 2013, we've shipped over 36 petabytes of flash, accelerating more than 3 exabytes of total storage.

We are one of the first major storage vendors to offer an enterprise-class flash array, and our flash portfolio is rounded out with a full range of hybrid storage and server flash options. More choices give you more degrees of freedom to choose the right solution at the right price. These solutions offer great performance plus the enterprise-class reliability and feature set you expect from NetApp.

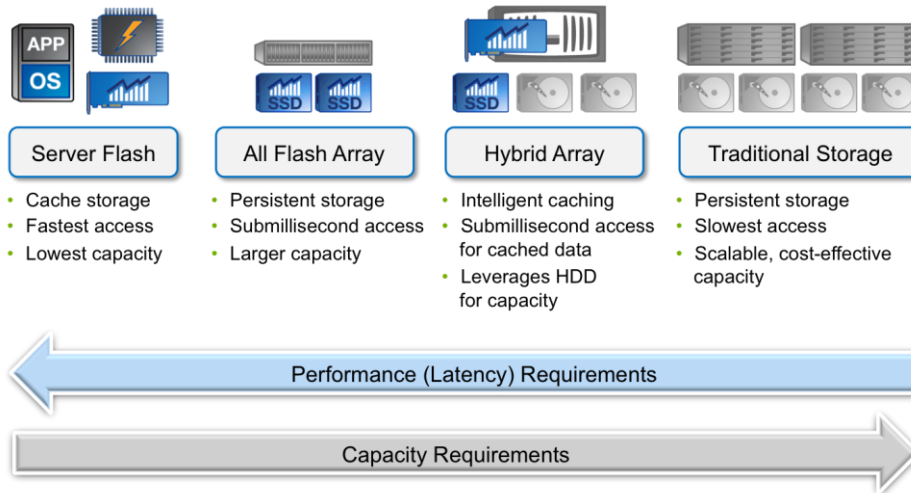
This white paper takes a detailed look at available enterprise flash options to help you understand which solutions make the most sense based on your business requirements. An introduction to the full NetApp® flash portfolio is included.

2 Deploying Flash for Enterprise IT

There are a number of options for deploying flash in the data center:

- **Hybrid storage solutions** combine the performance of flash with the capacity of HDD by targeting hot data to flash using either migration or caching.
- **Server flash solutions** may provide persistent solid state storage or cache data from HDD storage onto flash devices installed in servers, delivering extremely low latency for data accessed from cache.
- **All-flash arrays** provide maximum performance and a high level of consistency for business-critical applications.

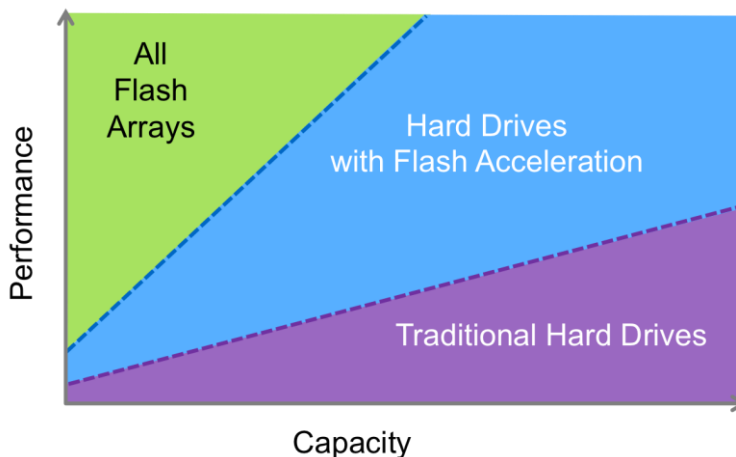
Figure 2) Enterprise storage options.



None of these solutions by itself can address all data center storage needs. The solution you ultimately choose could depend on a variety of factors including the size of your active dataset, how quickly the active dataset changes, and the number of applications that require acceleration.

NetApp performed a detailed analysis to determine the most cost-effective options across a wide range of performance/capacity combinations to reveal where each solution is most appropriate.

Figure 3) Performance/capacity analysis shows the best fit for different flash technologies.



All-flash arrays are most cost effective where capacity requirements are low and performance requirements are high. Hybrid configurations (including server flash) dominate the middle of the chart with performance increasing as HDD and flash capacity are added.

The dividing line between all-flash arrays and accelerated hybrid configurations shown in Figure 3 is moving to the right as flash technology becomes less expensive and incorporates storage efficiency technologies such as deduplication and compression to decrease overall flash consumption. As this happens, all-flash configurations will become appropriate for a broader range of requirements.

Many workloads currently fall on the border between the all-flash and hybrid zones. These are likely to move into the all-flash zone as the cost for flash capacity falls.

Workload requirements are also changing. The performance requirement of structured workloads is increasing while the capacity requirement remains relatively flat—making these workloads ripe for a change in media. For many unstructured workloads, capacity requirements are increasing faster than the drop in media prices, so hybrid solutions remain a better choice for now.

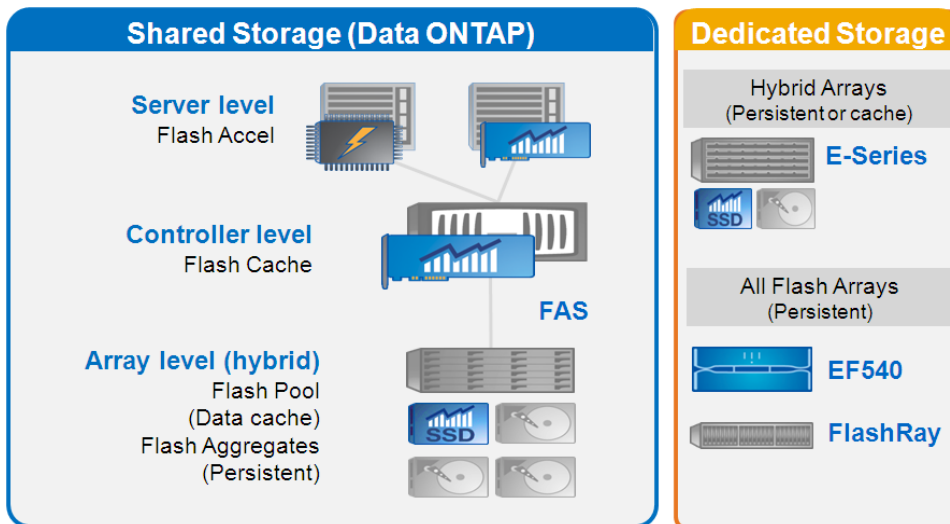
3 A Complete Flash Portfolio for Enterprise Needs

NetApp provides a complete flash portfolio so you can choose the solutions that deliver the performance and capacity you need—and tailor those solutions to the exact requirements of your applications. Flash has a role to play across the entire infrastructure stack; that’s why our portfolio includes hybrid storage, server flash, and all-flash array options.

- **Hybrid flash options** are available for both NetApp FAS and NetApp E-Series platforms. NetApp Flash Cache™ and Flash Pool™ intelligent caching for FAS and V-Series systems and SANtricity® SSD Cache for E-Series use flash resources to automatically cache hot data without complex configuration or ongoing management.
- **NetApp server flash** combines our Flash Accel™ software with your choice of flash hardware on servers to provide a persistent cache for applications requiring the lowest possible latency.
- **NetApp flash arrays** provide a dedicated, high-performance platform that delivers performance and reliability combined with enterprise features.

All of these options can reduce the total number of disk spindles required for a given workload. With hybrid storage and server flash you may also be able to deploy capacity-optimized disk drives rather than performance-optimized drives for further savings.

Figure 4) The NetApp flash portfolio includes FAS, E-Series and FlashRay platforms and encompasses the infrastructure stack from servers to disk arrays.



We recognize that a buying decision can be complex. Although performance, capacity, and cost provide a good starting point, in the end other factors such as availability, data protection, and the ability to provide enterprise-class service and support may be just as important. All NetApp solutions offer not only the performance, capacity, and scale you expect; they also deliver proven availability and the enterprise features to streamline your IT operations and help your business succeed—all backed by our worldwide support and services organizations.

Evaluating Flash Solutions

When evaluating flash solutions, there are four decision levers you must take into account. The more accurately you can estimate the values for these criteria, the better you can zero in on the best solution for your needs.

- **Performance.** What are your read and write performance requirements in terms of IOPS, and what is your latency target? How critical is consistent latency for every operation as opposed to meeting an average latency target?
- **Scale.** How much total capacity do you need and how large is your typical working set (the active data) that needs to be stored in flash?
- **Budget and space.** What are your cost-per-gigabyte (\$/GB) and cost-per-IOPS (\$/IOPS) price targets? How critical are space, power, and cooling in your decision? (A lower-cost solution that requires a costly data center upgrade is no bargain.)
- **Data management and protection.** What capabilities do you need for high availability, data protection, and disaster recovery? What are your support requirements?

Customer profiles throughout this paper explore these levers and how they lead to particular buying decisions.

4 Hybrid Storage Solutions

NetApp has been leading the deployment of hybrid storage since the introduction of the first Flash Cache product in 2009. We put a lot of time and energy into understanding the challenges that hybrid storage must address to architect solutions that deliver maximum application performance while minimizing cost and complexity.

From a performance perspective, the primary goal of hybrid storage is to shift as much random I/O as possible to flash to increase IOPS and reduce average latency to users and applications. We've found that a cache-based approach—in which hot data is automatically promoted to flash media—provides significant performance improvements for a wide range of applications without the requirement for setup, tuning or ongoing management. We use this approach with our Virtual Storage Tier (VST) technologies for the NetApp Data ONTAP® operating system as well as the SSD Cache available with NetApp SANtricity for the NetApp E-Series.

4.1 Data ONTAP and the Virtual Storage Tier

NetApp Virtual Storage Tier (VST) technologies used with NetApp Data ONTAP offer real-time promotion of data; hot data enters the cache immediately so there's no lag in performance acceleration and there's no need to set policies, thresholds, or time windows, making deployment and management simple. VST is fully integrated with the unified storage architecture of Data ONTAP, so it works with any storage protocol and the full suite of NetApp storage efficiency and integrated data protection capabilities to achieve optimal results.

Flash Cache

Flash Cache accelerates random read performance for all application workloads on a controller. It is well suited for technical, web, and business applications; file services; and VDI, making it an ideal choice for mixed storage workloads.

For example, a file services workload was originally being handled by a FAS6210 with 240 high-performance HDDs (600GB, 10K RPM). A 1TB Flash Cache was added and the existing HDDs were replaced with 168 high-capacity HDDs (1TB, 7.2K RPM). This resulted in:

- 34% lower cost per TB
- 40% lower cost per IOPS
- 17% more capacity
- 40% less power
- 28% more IOPS
- 18% reduction in average response time

Flash Pool

NetApp Flash Pool accelerates random read and write performance for specific workloads and is particularly well suited to OLTP applications. From the perspective of an HDD, the most "expensive" activities are random block reads and overwrites. Flash Pool off-loads these operations to SSDs while HDDs satisfy your capacity requirement. Caching writes prevents short-lived writes from being written to HDD.

A comparative before-and-after study using an OLTP workload illustrates the impact of Flash Pool. The 240 high-performance drives (600GB, 10K RPM) on a FAS6210 were replaced with 216 high-capacity disks (1TB, 7.2K RPM) and 24 SSDs (100GB). The updated configuration delivered the same IOPS with:

- 46% lower cost per TB
- 18% lower cost per IOPS
- 50% more capacity
- 26% less power consumption
- Significant improvement in response times

4.2 E-Series Solutions

NetApp E-Series storage is purpose-built for speed. E-Series systems are most often dedicated for use by high-performance applications such as real-time analytics, OLTP, and streaming video. The NetApp SANtricity SSD Cache is designed to optimize performance for these and similar application workloads.

E-Series technology is also widely sold through OEMs. Partners such as Teradata use an approach to hybrid storage that is optimized to the unique needs of particular applications.

SANtricity SSD Cache

Like Flash Pool, SANtricity SSD Cache automatically caches data blocks on SSD in real time, without the need for policies or scheduled data migration. The minimum cache is a single SSD drive, and SSD Cache can be shared by any or all volumes on an E-Series system.

SANtricity SSD Cache offers design optimizations to allow greater flexibility for high-performance applications including tunable block size, which we've shown to increase the speed with which the cache is populated. Write behavior can be set to match the characteristics of your application.

Workloads tested with SANtricity SSD Cache show up to a 700% increase in IOPS over the same arrays without cache.

Migration-Based Hybrid Storage

The key to hybrid storage is to make sure that active datasets are on flash. Although a caching approach works well for a broad set of applications, actively migrating data from HDD to SSD works well in situations in which application behavior is fully understood, particularly when the working set is so large that it would exceed the size of the cache.

For example, [Teradata](#), a leading NetApp OEM partner, uses this approach for its data warehouse hardware platforms. A Teradata data warehouse knows exactly how often each data block is accessed and hence where it is best stored. [Teradata Virtual Storage software](#) monitors how often data blocks are accessed and automatically migrates frequently accessed data to SSD.

The way that data is used in data warehouse systems is surprisingly similar across customers and industries. By monitoring hundreds of systems, Teradata found that data access follows the classic 80:20 rule—20% of the data is accessed 80% of the time. By committing 20% of the capacity in each array to SSDs, users achieve up to a 60-times improvement in analytics operations.

Optimal Media: Hybrid Storage

Hybrid storage is the most cost-effective choice for a broad range of capacity and performance requirements. (See Figure 3.)

- **Performance.** 40K–200K IOPS (combined read and write), 3–5ms latency.
- **Scale.** 40TB–1PB capacity, <2TB working set. Multiple applications need to share a common storage infrastructure and workloads are unpredictable.
- **Budget and space.** Storage footprint reduced by decreasing the number of spindles to achieve a given level of performance or by replacing performance optimized disk with capacity optimized disk.
- **Data management and protection.** Mature data management and protection options.

5 Server Flash

NetApp **Flash Accel™** software allows hot data to be cached on servers as close as possible to the application, minimizing latency for data in cache while providing a level of persistence, durability, and cache coherency that other server flash solutions lack. Our approach combines the benefits of fast local storage with the convenience of consolidated primary storage for data management and protection. It works particularly well when the application to be accelerated has a modest active dataset—so it fits in flash on the server—and the active dataset isn't changing rapidly.

Flash Accel works with any flash devices on your servers, protecting existing investments. It increases throughput by up to 80%, and reduces transaction latency by up to 90%. NetApp recommends deploying Flash Accel in conjunction with Flash Cache or Flash Pool on primary storage for best results.

We compared the performance of the same configuration with and without Flash Accel using JetStress, which simulates the disk I/O load created by Microsoft® Exchange. The addition of Flash Accel resulted in approximately a 77% improvement in I/O performance for both reads and writes.

In another example, Flash Accel was used to optimize an actual Microsoft Exchange configuration. The original configuration used a FAS3240 with 512GB Flash Cache and 72 high-capacity disk drives (1TB, 7.2K RPM). The updated configuration added Flash Accel and a PCIe flash card to the server and reduced the disk count to 48, a 33% reduction in total spindles. Performance remained the same while total storage cost was reduced by 15% and memory and CPU utilization on the FAS3240 was reduced by 10–15%.

Optimal Media: Hybrid Storage plus Server Flash

Server flash is a good addition to a hybrid storage configuration in situations in which a critical performance objective is isolated to a specific application running on one or a few servers. It works best in situations in which the active dataset—the working set—is relatively small (to fit in flash on the server) and situations in which consistent performance for every I/O is not a requirement.

- **Performance.** 40K–400K IOPS (combined read and write), 1–2ms latency.
- **Scale.** 40TB–1PB capacity, <2TB working set. Multiple applications and workloads share a common storage infrastructure. One or a few applications require additional latency acceleration.
- **Budget and space.** Storage footprint reduced by decreasing the number of spindles to achieve a given level of performance or by replacing performance optimized disk with capacity optimized disk. Additional efficiency in storage controller utilization.
- **Data management and protection.** Data management and protection provided by back-end storage arrays.

6 All-Flash Arrays

NetApp began shipping its first all-flash array—the EF540—in 2012. The EF540 is designed to deliver superior performance, efficiency, and availability in a platform that leverages the knowledge and experience gained from a legacy of over 500,000 deployed storage systems. A new, purpose-built flash array (FlashRay) is also in the works and will be available in 2014.

6.1 The EF540

The EF540 is the industry's first all-flash array that delivers flash performance plus enterprise-grade reliability and features. These capabilities can improve the speed of your business, as well as the overall effectiveness of your IT operations.

The EF540 is intended for performance-driven applications with submillisecond latency requirements. With the EF540, your business-critical applications deliver results faster—achieving response times only possible with all-flash storage. The combination of high IOPS and ultralow latency makes the EF540 flash array a great system for database-driven applications.

A single system delivers up to 19.2TB of raw capacity and over 300,000 IOPS of predictable, repeatable and sustainable performance with extremely low latency. This is equivalent to the transactional performance of over 1,000 traditional 15K hard disks drives but uses 95% less space, power, and cooling. It's also worth noting that the EF540 can deliver up to 6 gigabytes/sec throughput for large-block I/O or bandwidth-oriented workloads.

The EF540 offers full enterprise-class reliability and availability features. Fully redundant I/O paths and automated failover, which—perhaps surprisingly—are not available on some flash arrays, are standard. All management tasks can be performed while the EF540 remains online with complete read/write data access. Extensive monitoring of diagnostic data provides comprehensive fault isolation, simplifying analysis of unanticipated events for timely problem resolution. Proactive repair can help get the system back to optimal performance in minimum time.

Advanced data protection options include high-efficiency snapshots for local protection and asynchronous or synchronous replication for disaster recovery.

6.2 FlashRay

In February 2013 NetApp announced the FlashRay™ storage array, a purpose-built all-flash storage architecture that features scale-out and combines consistent low-latency performance, high availability, and integrated data protection with enterprise storage efficiency features such as inline deduplication and compression. FlashRay will begin beta testing in mid-2013 and will be generally available in 2014.

Optimal Media: All-Flash Array

All-flash arrays are the best choice when all transactions must be completed with consistently low latency. It is a particularly good fit when the performance requirement is high and the capacity requirement is modest.

- **Performance.** 50K–400K read IOPS, <1ms latency with high consistency.
- **Scale.** 15TB–40TB capacity, <40TB working set. Heterogeneous and unpredictable workload. Storage needed across numerous geographic locations.
- **Budget and space.** Reduced \$/IOPS and data center space by reducing the number and types of spindles needed to support an application.
- **Data management and protection.** Because there are many new vendors in this space, data management and protection options may vary widely and be a key differentiator in situations in which performance and cost are comparable.

7 Conclusion

Flash technology makes applications both faster and more responsive, so the processes your business relies on finish sooner and your IT efficiency increases. As a leading storage vendor, NetApp has the experience with flash technology, the broad flash portfolio, and the enterprise capabilities required to address a wide variety of flash needs across your IT infrastructure, including all-flash arrays, a full selection of hybrid flash solutions, and server flash.

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®



www.netapp.com

© 2013 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, Data ONTAP, Flash Accel, Flash Cache, Flash Pool, FlashRay, and SANtricity are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Microsoft is a registered trademark of Microsoft Corporation. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. WP-7182-0413