

Compliments of  NetApp®

cStor

NetApp 2nd Special Edition

Flash Storage

FOR
DUMMIES[®]
A Wiley Brand

Learn to:

- Optimize storage performance
- Leverage flash in hybrid array storage with intelligent caching
- Identify workloads well suited for all-flash storage
- Reduce the footprint of your storage infrastructure

Lawrence C. Miller, CISSP



About NetApp

NetApp creates innovative storage and data management solutions that deliver outstanding cost efficiency and accelerate business breakthroughs. Discover NetApp's passion for helping companies around the world go further, faster at netapp.com

About cStor

cStor designs smart data center and cloud technology solutions that streamline operations, reduce costs, improve security, create competitive advantage and move business forward faster. Learn more about our passion for using best-of-breed technologies to solve IT challenges of all shapes and sizes at cstor.com.

Flash Storage

FOR
DUMMIES[®]
A Wiley Brand

NetApp 2nd Special Edition

by Lawrence C. Miller, CISSP

FOR
DUMMIES[®]
A Wiley Brand

Flash Storage For Dummies®, NetApp 2nd Special Edition

Published by
John Wiley & Sons, Inc.
111 River St.
Hoboken, NJ 07030-5774
www.wiley.com

Copyright © 2014 by John Wiley & Sons, Inc., Hoboken, New Jersey

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, the Wiley logo, For Dummies, the Dummies Man logo, A Reference for the Rest of Us!, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. NetApp, the NetApp logo, Go further, faster, Data ONTAP, Flash Cache, and Flash Pool are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact info@dummies.biz, or visit www.wiley.com/go/custompub. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN 978-1-118-99994-3 (pbk); ISBN 978-1-119-02958-8 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

These materials are © 2014 John Wiley & Sons, Inc.
Any dissemination, distribution, or unauthorized use is strictly prohibited.

Contents at a Glance



Introduction	1
Foolish Assumptions	1
About This Book	1
Icons Used in This Book.....	2
Where to Go from Here	3
Chapter 1: Flash Storage Technology — What and Why	5
Flash Concepts	5
SLC, MLC, and Some TLC	6
The Evolution of Flash Technology	8
Flash in the Enterprise	10
Chapter 2: Flash Storage Technology — Where and How	17
Hybrid versus All-Flash Approaches	18
Storage Array Flash as Cache.....	19
SSD as Cache	23
SSD Replacement for HDD	24
Chapter 3: Recognizing the Benefits of Flash Storage Technology	27
Performance	27
Reliability	28
Efficiency.....	30
Cost.....	34
Chapter 4: Ten (Okay, Five) Strategies for Using Flash in Your Next Storage Project.....	39

Publisher's Acknowledgments

Some of the people who helped bring this book to market include the following:

Senior Project Editor: Zoë Wykes

Editorial Manager: Rev Mengle

Acquisitions Editor: Kyle Looper

**Senior Business Development
Representative:** Karen L. Hattan

**Custom Publishing Project
Specialist:** Michael Sullivan

Project Coordinator:

Melissa Cossell

Special Help from NetApp:

Larry Freeman, Paul Feresten,
Darrell Suggs, Duncan Moore,
Cesar Orosco, Rip Wilson

Introduction



Flash technology is making quite a splash in the storage industry. Offering superior speed and reliability when compared to traditional hard disk drives, flash storage is a flexible and increasingly cost-effective technology that can be used to optimize enterprise storage environments running mission-critical, I/O-intensive applications.

Flash Storage For Dummies, NetApp 2nd Special Edition, explores the many uses and benefits of flash storage technology in the enterprise. From flash-accelerated storage to all-flash arrays, flash technology improves performance and increases reliability in storage infrastructures. It also reduces energy and real-estate costs in the data center.

Foolish Assumptions

It's been said that most assumptions have outlived their usefulness, but I'll assume a few things nonetheless. Mainly, I assume that you're a technical manager or strategist with some understanding of storage technologies — such as network-attached storage (NAS) and storage area networks (SANs). As such, this book is written primarily for technical readers who are evaluating new storage technologies.

About This Book

This book contains volumes of information that rival the U.S. Congressional Record or the complete

Encyclopedia Britannica, conveniently distilled into four short chapters that are chock-full of just the information you need. Here's a brief look at what awaits you in the pages ahead.

Chapter 1: Flash Storage Technology — What and Why. I start by taking a look at flash technology: what it is, how it evolved, and why you should consider it for your storage infrastructure.

Chapter 2: Flash Storage Technology — Where and How. Here, you find out about NetApp's flash storage solutions: NetApp Flash Cache, Flash Pool, all-flash FAS, and the EF-Series.

Chapter 3: Recognizing the Benefits of Flash Storage Technology. This chapter tells you how flash storage can benefit your organization.

Chapter 4: Ten (Okay, Five) Strategies for Using Flash in Your Next Storage Project. Finally, in that classic *For Dummies* style, I give you some useful tips on how to use flash in your storage environment.

Icons Used in This Book

Throughout this book, you occasionally see icons that call attention to important information that is particularly worth noting. You won't find any winking smiley faces or other cute little emoticons, but you'll definitely want to take note. Here's what to expect.



This icon points out information that may well be worth committing to your nonvolatile memory, your gray matter, or your noggin — along with anniversaries and birthdays.



If you're a long-suffering insomniac or vying to be the life of a World of Warcraft party, take note. This icon explains the jargon beneath the jargon and is the stuff legends — well, at least nerds — are made of.



Thank you for reading, hope you enjoy the book, please take care of your writers! Seriously, this icon points out helpful suggestions and useful nuggets of information.

Where to Go from Here

With our apologies to Lewis Carroll, Alice, and the Cheshire Cat:

“Would you tell me, please, which way I ought to go from here?”

“That depends a good deal on where you want to get to,” said the Cat — er, the Dummies Man.

“I don't much care where . . .,” said Alice.

“Then it doesn't matter which way you go!”

That's certainly true of *Flash Storage For Dummies*, NetApp 2nd Special Edition, which, like *Alice in Wonderland*, is destined to become a timeless classic.

If you don't know where you're going, any chapter will get you there — but Chapter 1 might be a good place to start. However, if you see a particular topic that piques your interest, feel free to jump ahead to that chapter. Read this book in any order that suits you (though I don't recommend upside down or backward).

I promise that you won't get lost falling down the rabbit hole.

Chapter 1

Flash Storage Technology — What and Why

In This Chapter

- ▶ Understanding how flash works
 - ▶ Taking flash to new levels
 - ▶ Studying up on the history of flash
 - ▶ Saying goodbye to HSM, ILM, and AST
-

In this chapter, you learn how flash works, how it has evolved as a viable — even preferred — storage technology for consumers and enterprises, and how it enables new data and storage management strategies.

Flash Concepts

Flash storage technology operates very differently from traditional hard disk drives (HDDs) that store data on rotating aluminum platters that are magnetically coated surfaces.



Flash technology stores information in an array of columns and rows with a memory cell at each intersection. Each cell has two transistors —

a floating gate and a control gate — separated by a thin oxide layer. Each cell contains a single bit of information; applying an electrical charge to the floating gate transistor determines whether the cell represents a 1 or a 0. This electrical charge is how flash got its name, like the flash of a camera, each time data is stored. Multi-level cells (MLC) and triple-level cells (TLC), discussed in the next section, can store more than one bit per cell by applying varying levels of electrical charge to their floating gates.

SLC, MLC, and Some TLC

Flash technology has many advantages, but it also has two major weaknesses today:

- ✓ Like all forms of electrically erasable programmable read-only memory (EEPROM), those little transistors inside the memory module don't like being flashed. After multitudes of electrical erasures, the transistor gates begin to break down, and eventually they fail.
- ✓ Because you always need to erase flash before writing new data to it, in some cases flash isn't very fast. In fact, flash-based solid state drives (SSDs) can be slower than high-performance HDDs under some conditions, such as large sequential write operations. However, for other operations, such as random reads, SSDs are incredibly faster than HDDs.



Thefreedictionary.com defines EEPROM (Electrically Erasable Programmable ROM) as:

A rewritable memory chip that holds its content without power. EEPROMs are bit or byte addressable at the write level, which means either the bit or byte must be erased before it can be re-written. In flash memory, which evolved from EEPROMs and is almost identical in architecture, an entire block of bytes must be erased before writing. In addition, EEPROMs are typically used on circuit boards to store small amounts of instructions and data, whereas flash memory modules hold gigabytes of data for digital camera storage and hard disk replacements

To address these weaknesses, several new technologies are being developed that could position flash technology to someday replace HDDs altogether.

For example, single-level cell (SLC) flash has traditionally been used in enterprise SSDs. SLC flash is ten times more durable than MLC flash, albeit at a much higher cost.

To get the durability of SLC flash in MLC flash without the high cost, some storage companies have started to offer MLC flash with sophisticated wear-leveling and bad-block management algorithms. These innovations enable lower-cost MLC flash to be used in enterprise SSDs without sacrificing reliability.

TLC flash provides dramatically higher flash capacities and still lower costs, but it also has slower access times and lower reliability than even MLC flash. However, as data storage pattern algorithms become even more sophisticated and durability increases, TLC flash is likely to make its way into enterprise SSDs.

The Evolution of Flash Technology

Every once in a while a technological innovation comes along that revolutionizes an entire industry. Flash storage is one such innovation — a disruptive technology — that is fundamentally changing the storage industry.



A disruptive technology is an innovation that displaces an earlier technology in an existing market and value network.

The following sections highlight the evolution of flash technology.

Read-only memory

For all its flash (sorry), flash technology traces its roots to a very humble beginning — read-only memory (ROM). In the early days of integrated circuits, circa the 1970s, computer microcode was permanently stored on ROM chips. As their name implies, ROM chips always held the same data, or programming instructions, and could not be erased or rewritten. ROM worked great until a code upgrade was needed, which required the ROM chipsets to be replaced with newer ROM chipsets containing the upgraded code.

Erasable programmable read-only memory

Erasable programmable read-only memory (EPROM) was the next evolution of ROM. EPROM overcame the permanent “write once” limitation of ROM, thus providing a major breakthrough for the computing and storage industries.

EPROM had a clear window over its silicon chip and a little sticker over the window, usually denoting the revision (or version) of the microcode on the chip. When it was time to upgrade the code, you pulled the sticker off and put the chip into a little box with an ultraviolet light — a microwave oven for computer engineers! After about an hour under the light, the silicon chip was erased and could be recoded using a special EPROM programmer.

EPROMs allowed you to reuse expensive memory chips, but they were still too cumbersome. The third evolution of memory was EEPROM. With EEPROM, a chip could be erased by simply applying an electrical charge across its memory cells. In fact, EEPROMs could be erased and reprogrammed without even removing them from the computer system's motherboard.

EEPROMs were useful for more than just storing microcode; they could also be used as solid state storage devices like their cousins, random-access memory (RAM) chips. In fact, because they could store, erase, and store data again, EEPROMs were categorized as a type of nonvolatile random access memory (NVRAM).

EEPROMs took a huge leap forward as consumer electronics products, including digital music players, cameras, and video recorders — all of which needed substantial storage capacity — became popular. But in the early 1990s, the only viable storage technology for these consumer electronics products was a 1-inch microdisk drive, which proved unreliable and expensive.

Compact Flash

In 1994, Sandisk announced its Compact Flash module. Compact Flash brought innovation to the consumer

electronics market in a standard packaged module that was a plug-and-play replacement for the aforementioned microdisk drives. Key to the success of Compact Flash was the ability to erase and rerecord single memory cells, unlike earlier EEPROMs that had to be completely erased before any new data could be stored.

Thus a new era of storage was born and the rest, as they say, is history! As the demand for consumer electronics took off, the cost of flash technology plummeted to the point where storage vendors took notice, and many IT professionals are now asking “Could flash completely replace traditional hard drives?”

Flash in the Enterprise

As enterprise data storage requirements have grown at unprecedented rates, in terms of both capacity and performance, the storage industry has struggled to deliver innovative technologies that effectively address today’s enterprise storage challenges.

One of the big ways that flash storage is affecting enterprise data centers is in storage tiering, a concept that has come and gone — again and again. In the early 1980s, storage tiering was known as hierarchical storage management (HSM). In the 1990s, HSM became information lifecycle management (ILM). Now, in its latest incarnation, storage tiering is referred to as automated storage tiering (AST). Flash technology has breathed new life into storage tiering, using new implementations that modernize the whole concept of tiering.

Although the name has evolved, the fundamental principle is the same: Storage tiering is based on the outdated and operationally inefficient methodology of

moving data from high-performing, very-expensive media to high-capacity, less-expensive media, with the goal of maximizing the utilization of the storage infrastructure. But whatever its current name, storage tiering is a project that IT organizations never seem to get done.

The problem is that most storage tiering solutions available today are overly complex, cumbersome to implement, and based on an antiquated methodology (see Figure 1-1).

- Monitor I/O behavior for groups of disk blocks (chunks)
- Compare access frequency against average for each tier
- Promote or demote chunks to suitable tier

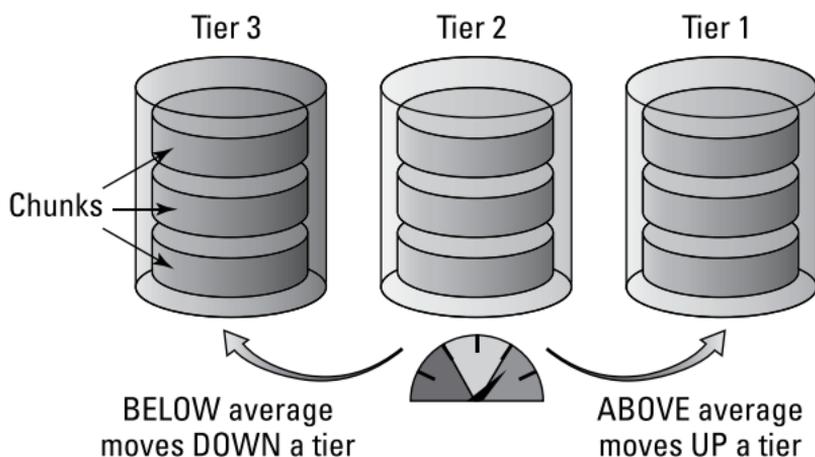


Figure 1-1: How storage tiering works.

The basic premise of storage tiering is about storing data in the right place and at the right time and price to support the enterprise and to utilize the storage infrastructure most efficiently.

Three basic assumptions underlie the premise of storage tiering:

- ✔ The value of data decreases over time; according to some estimates, data not accessed within 90 days will almost never be accessed again.
- ✔ It is estimated that less than 20 percent of all data needs to be on high-performance (and therefore expensive) media.
- ✔ It's a reasonable strategy to move as much data as possible to high-capacity serial ATA (SATA) drives as soon as possible as a way to better utilize the storage infrastructure in the data center.

So how do today's automated tiering solutions attempt to address these three points? Let's start by focusing on the word "automated"; it sounds right, and it's very appealing — but does it work? Before IT departments can "automate" anything, a tremendous amount of work needs to be done. The storage architect or administrator has to collect, analyze, and design the correct workflows so that the system can automate them — in other words, the storage architect has to do the heavy lifting.

For example, consider the following questions that need to be answered in order to architect an appropriate solution:

- ✔ How many tiers of storage does my environment need? Some storage vendors offer as many as nine tiers of storage based on type of drive, rotational speed, and RAID (redundant array of independent disks) level.
- ✔ How big should tier 1 be? How big should tiers 2, 3, 4, and so on be?

- ✔ How do I determine what data is “hot,” “warm,” or “cold” (discussed in Chapter 2)? Some implementations of automated storage tiering require a good understanding of application workloads, additional software, and detailed planning and sizing of the different tiers of storage.
- ✔ What kind of data should go in tier 1? How do I classify my data?
- ✔ How long does it take auto-tiering software to migrate data to another tier? In some instances, it takes 3 days to promote data and 12 days to demote it.
- ✔ When is critical data promoted to tier 1? Keep in mind that data migrations or relocations can impact system performance; depending on the vendor, it could take hours to days.
- ✔ When is cold data moved to tiers 2 and 3?
- ✔ Is the data migration process manual, automatic, or scheduled?
- ✔ How granular is the data migration? Do I need to move a whole logical unit number (LUN)? Or a sub-LUN?
- ✔ How do I know if I have the right data migration policies, thresholds, or time windows for data movement? Ongoing monitoring or calibration will be required.
- ✔ Can I use data-efficiency features like deduplication and thin provisioning in my tier 1 storage layer?
- ✔ Do I need different tiering solutions for NAS and for SANs?
- ✔ How many new tools and management end points will storage tiering add to my environment?

- ✓ How much will the storage tiering cost? How many licenses will I have to purchase?

Architecting the right solution depends on making sure that these questions get answered correctly and that you've collected and analyzed the correct data. In the end, even if your data and analysis are correct, the actual implementation of the solution may be too complex. The success or failure of storage tiering requires painstaking work up front and constant feeding and care in the operations phase. It assumes predictable workloads, and there is little flexibility.

A far less complex and much better approach is to start with unified storage — a single platform that combines relevant media types, array functionality, and all protocols. This solution allows you to use software to build a virtual storage tier (VST) strategy using flash to optimize performance where you need it and to reduce storage costs where you don't.

VST is a self-managing, data-driven service layer for storage infrastructure (see Figure 1-2). VST provides real-time assessment of workload priorities and optimizes I/O requests for cost and performance without the need for complex data classification and movement.

VST promotes hot data to flash storage without the data movement or migration overhead associated with other approaches to storage tiering. Whenever a read request is received for a block on a volume or LUN where VST is enabled, that block is automatically subject to promotion (4KB blocks are very granular, compared to other implementations). Promotion of a data block to the VST is not data migration, because the data block remains on the hard disk when a copy is made to the VST.

Virtual Storage Tier

Data-driven • Real-time • Self-managing

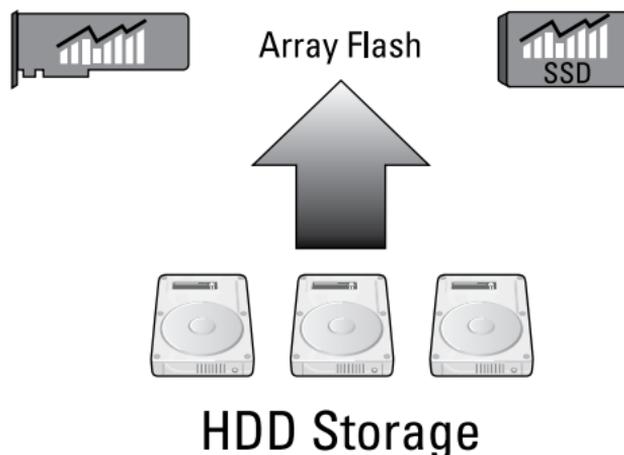


Figure 1-2: Virtual storage tiering with flash technology.

VST leverages key storage efficiency technologies, such as deduplication, volume cloning, thin provisioning, intelligent caching, and simplified management. You simply choose the default media tier you want for a volume or LUN, such as SATA, Fibre Channel (FC), or serial-attached SCSI (SAS). Hot data from the volume or LUN is automatically promoted on demand (application-driven) to flash-based media.

A key difference between VST and past storage tiering approaches is the innovative use of flash technology. In Chapter 2, you learn about NetApp's flash technology solutions and how they work together to deliver a complete VST strategy in the enterprise.

Chapter 2

Flash Storage Technology — Where and How

In This Chapter

- ▶ Choosing your flavor: hybrid or all flash
 - ▶ Replacing SAS with SATA and flash
 - ▶ Taking a deep dive into NetApp Flash Pool
 - ▶ Going all the way with all-flash storage
-

Flash-based storage improves I/O performance and efficiency for many applications, including databases, server and desktop virtualization, and cloud infrastructure. Chapter 1 explains what flash technology is and why it is a viable and innovative enterprise storage technology that enables solutions such as virtual storage tiering. In this chapter, you get acquainted with NetApp's hybrid and all-flash array solutions, as well as how NetApp uses flash for storage tiering to maximize the value of flash technology.

Hybrid versus All-Flash Approaches

Deploying flash technology at the storage array level offers two basic options:

- ✓ Hybrid
- ✓ All flash

Hybrid storage arrays

Hybrid storage arrays combine the performance of flash with the capacity of HDD by targeting hot data to flash using either migration or caching techniques. 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.

NetApp provides hybrid array options for both its Data ONTAP platform (Flash Cache, Flash Pool) and its SANtricity platform (SSD Cache). In these hybrid solutions, a relatively small amount of flash technology, typically 1–2 percent of total storage capacity, can yield dramatic efficiency and performance benefits.

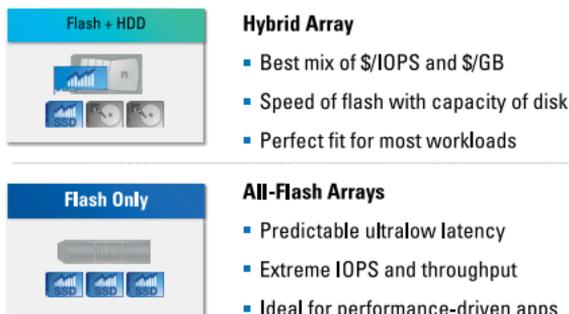
All-flash arrays

All-flash arrays provide maximum performance, fast response times, and a high level of consistency for business-critical applications. All-flash arrays are the best choice when all I/O operations must be completed with consistently low latency, typically below 1 millisecond (ms). These are a particularly good fit when the performance requirement is high and the capacity requirement is modest.

NetApp provides all-flash array solutions for both Data ONTAP and SANtricity platforms, covering a full range of performance and data management capabilities.



The choice between a hybrid array and an all-flash array will be based on the performance and efficiency requirements of the application workloads for which they are deployed. Those applications that demand the highest and most consistently predictable performance are the best candidates for all-flash deployment. See Figure 2-1.



Source: NetApp

Figure 2-1: Criteria for choosing hybrid versus all-flash arrays.

Storage Array Flash as Cache

Networked storage systems often struggle to meet enterprise performance demands. Adding capacity to these systems is easy enough — just add more disk drives. But disk drives aren't getting any faster — 10K or 15K RPM for SAS and 7.2K RPM for SATA. As a result,

large numbers of small-capacity disk drives are commonly used to deliver the I/O throughput demanded by many workloads. This approach wastes storage capacity, rack space, electricity — and money!



Short stroking is a technique that is sometimes used to improve IOPS by limiting the distance that the heads of a disk drive must travel to read data from any point on the disk drive. This reduces the average seek time but also reduces capacity.

NetApp Flash Cache improves performance for workloads that are random read-intensive without adding more high-performance disk drives. Flash Cache is a flash-based PCI-e expansion card for NetApp FAS storage controllers. Flash Cache is designed to improve workloads, such as

- ✓ Business processing
- ✓ IT infrastructure (file services)
- ✓ Collaboration (e-mail and Microsoft SharePoint)
- ✓ Engineering applications (software development and electronic design automation)

By providing an additional caching layer and special filtering software, the system is able to have a higher number of operations serviced from a low-latency medium. This has two effects:

- ✓ Higher throughput than a similar system without Flash Cache
- ✓ Fewer operations serviced by the disk subsystem, which can also present lower latencies of the from-disk IOPS

On the surface, all flash looks pretty much alike. Beneath the surface, however, design decisions can make all the difference. With Flash Cache, NetApp has forged a different path than other enterprise flash providers by focusing on software intelligence to maximize the efficiency of flash.

NetApp systems running Flash Cache have exhibited throughput in excess of 250,000 IOPS with latency below 1 millisecond. Latency is typically 10ms or higher when accessing data from disk drives. Flash Cache modules reduce latency by a factor of 10 or more compared to disk drives when a cache hit occurs (see Figure 2-2).

Flash Cache has also demonstrated that 15K RPM high-performance Fibre Channel and SAS hard disk drives can be replaced with fewer, lower-cost 7.2K RPM high-capacity SATA HDDs without sacrificing performance.

NetApp's latest Flash Cache modules (Gen4) contain 2 terabyte (TB) of flash. A maximum of eight modules can be installed in FAS storage controllers for a maximum per-controller flash capacity of 8TB. These modules contain MLC flash (see Chapter 1) managed by an onboard field-programmable gate array (FPGA). The FPGA controls all communication between the system main memory and the flash memory located on the Flash Cache board. Designed for speed, the FPGA employs many techniques to improve performance, such as:

- ✔ Intelligently interleaving writes throughout multiple write queues results in balanced flash erase, write, and read cycles. Because each flash write requires a flash erasure and erase cycles are very slow, it's best to have as many cycles as possible running in parallel to maintain throughput.

- Supporting multiple memory interfaces — each interface goes several banks deep. When one flash bank on an interface is busy (for example, during an erase cycle), the FPGA can issue a command to another bank on the same interface. This prevents stalls when there are too many requests and not enough flash memory banks.

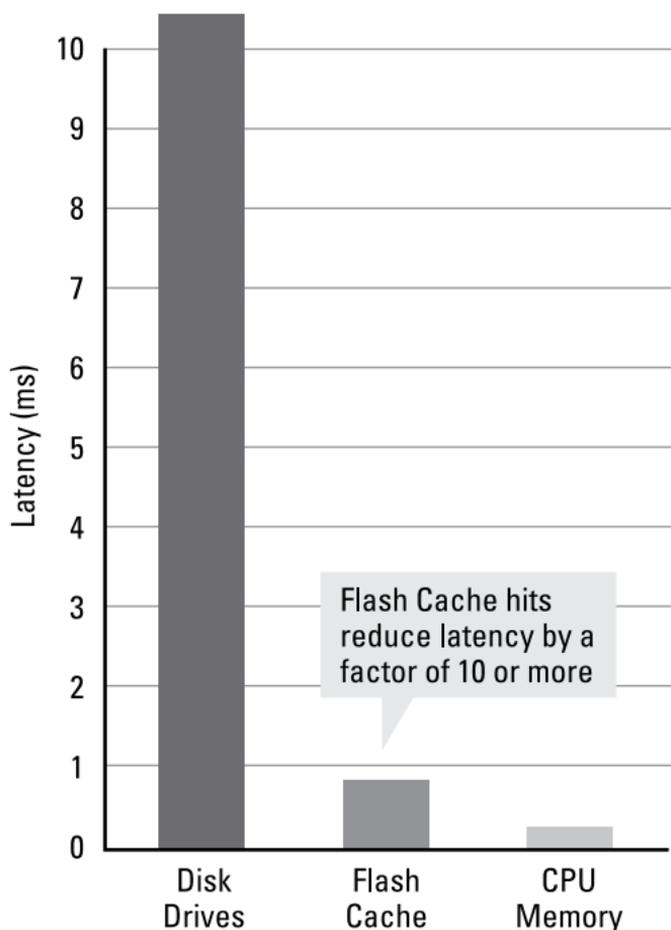


Figure 2-2: Latency reduction.

- ✓ Reading from flash cells in groups that are striped across multiple flash banks rather than reading them individually reduces read latency by more than 800 percent.

SSD as Cache

Combining SSDs and HDDs in a “hybrid” storage pool allows networked storage systems to take advantage of the low-latency and high-throughput benefits of SSDs, while maintaining the mass storage capacity of HDDs. NetApp offers its Flash Pool solution in all NetApp FAS and V-Series storage systems. On NetApp E-Series systems, similar functionality is provided in an “SSD cache.”

Using a combination of SSD and HDD within a single pool allows hot data to be tiered automatically between HDD and SSD for optimal performance (see Figure 2-3). Like Flash Cache (discussed earlier in this chapter), hot data is not moved to SSD; rather, a copy of the data is created in SSD. Copying rather than moving data has the following benefits:

- ✓ Copying data is a faster I/O process than moving it. (Moving it requires the system to copy the data and then delete it from its original location.)
- ✓ Once the hot data is ejected from the SSD, there is no need to rewrite the data back to HDD.

However, unlike Flash Cache, both reads and writes can be cached in Flash Pool SSDs. As discussed in Chapter 1, in some cases, like sequential write operations, HDDs actually outperform flash-based SSDs. For that reason, when a write operation occurs on a Flash Pool

enabled volume, logic is used to determine whether it would be faster to write to SSD and HDD together or just to HDD. Conversely, when a read operation occurs, logic determines whether data should be cached to SSD or read directly from HDD without caching. This determination is based primarily on whether the data pattern is random or sequential. In all cases, Flash Pool algorithms are optimized for speed.

Primary Storage

```

10101001110101001110110101110
00101101010101001110101001110
11010111000101101010101001101
10110101010100110010110101010
10011101010011101010011101101
01110001011010101010011101010
01110110101110001011010101010
01101101101010101001100101101
10101001110101001110110101110
00101101010101001110101001110
11010111000101101010101001101
10110101010100110010110101010
10011101010011101010011101101
01110001011010101010011101010
01110110101110001011010101010
01101101101010101001100101101
10101001110101001110110101110
00101101010101001110101001110
11010111000101101010101001101
10110101010100110010110101010
  
```

Objective:
Intelligently
place “hot”

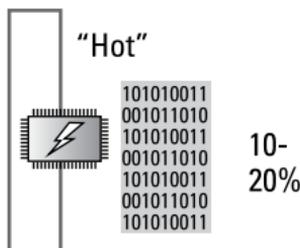


Figure 2-3: The data placement challenge: Make data fast when it’s hot and low cost when it’s cold.

SSD Replacement for HDD

Increasingly, enterprises are looking for ways to drive greater speed and responsiveness from the applications that they use in their key business operations — especially database systems. Because the performance

of these applications is tightly linked to time to market, revenue, productivity, and customer satisfaction, it is critical that they operate at maximum potential.

To achieve extreme performance, organizations may overprovision capacity in their networked storage systems, deploying partially filled disks that allow the array to deliver the required performance but also wasting costly resources including hardware, capacity, data center space, and power.

To eliminate overprovisioning and maximize return on investment from high-performance applications, many companies are now looking to all-flash storage systems. But as buyers consider these new systems, they are often challenged to find an offering that is enterprise class, highly reliable, feature rich, and cost effective.

NetApp all-flash arrays provide ultra-high performance, consistent submillisecond responsiveness, and enterprise features to drive greater speed, responsiveness, and value from the latency-sensitive applications that control key business operations. The NetApp EF-Series flash array is an enterprise-proven platform with a highly efficient, streamlined operating system that excels at I/O-intensive workloads where the focus is on extreme performance, ultra-low latency, and IOPS density. NetApp all-flash FAS systems combine low-latency performance with the feature-rich capabilities of Data ONTAP, including deduplication, compression, integrated data protection, multiprotocol support, scale-out, and nondisruptive operations.

Chapter 3

Recognizing the Benefits of Flash Storage Technology

.....

In This Chapter

- ▶ Improving storage performance with flash
 - ▶ Trusting your data to solid state technology
 - ▶ Optimizing storage infrastructures with SSD
 - ▶ Measuring the real cost of your storage
-

In this chapter, you find out how flash storage technology can benefit your organization, as well as how several companies in different industries have successfully deployed flash storage technology and are now enjoying its benefits.

Performance

Flash technology improves overall IT performance. Flash storage performs read and write operations in microseconds and provides tens of thousands to millions of IOPS. By comparison, hard disk drive (HDD) performance is rated in milliseconds and provides hundreds of IOPS (see Table 3-1).

Table 3-1 Flash versus HDD performance

	<i>Read/Write</i>	<i>IOPS</i>
Flash	Microseconds	Tens of thousands to millions
HDD	Milliseconds	Hundreds

And unlike HDDs, flash has no moving parts, so there are no performance hits due to seek time or rotational latency in the hardware. Having no moving parts in flash also means less wear and tear, which improves mechanical reliability (discussed in the next section).



Flash storage can improve response times by up to 90 percent and increases I/O throughput by up to 80 percent over HDDs.

See the case study at the end of the chapter, “NetApp EF-Series flash array accelerates Oracle database.”

Reliability

Reliability is, of course, one of the most important attributes of any data storage medium. You have to be confident that you are storing your company’s vital records and data in a reliable and consistent manner that ensures it will always be available when needed.

As mentioned earlier, flash storage has no moving parts. This design factor alone greatly improves reliability and results in a mean time between failures (MTBF) of approximately 2 million hours, compared to approximately 1.2 million hours for HDDs. See the sidebar “The future of the disk drive” to learn more about HDD design and reliability.

The future of the disk drive

One of the more lively discussions these days is about the future of the hard disk drive. HDDs have been around for more than half a century. The very first disk drive, the IBM 350 Disk Storage Unit, was composed of three basic elements: rotating aluminum platters, magnetically coated surfaces, and movable recording heads. Today, although the speed and capacity of HDDs have increased, they are still composed of these same three basic elements and function on the same operating principles.

Although the HDD has had an amazing run, someday it *will* be replaced. Precision motors and actuators are subject to mechanical wear and do not perform indefinitely. Spinning drive platters 24 hours a day but accessing data only occasionally is an inefficient use of electrical energy. Finally, while heat-assisted magnetic recording (HAMR) technology is expected to take HDD capacities to at least 100TB, at some point a capacity, reliability, and economic ceiling for HDDs will be reached.

The reason the HDD has survived so long is that nothing better has come along worthy of replacing it — until now. Enterprise SSDs use flash technology and emulate the look and feel of an HDD, but because enterprise SSDs are made entirely of silicon, they are not subject to the mechanical wear or the rotational latency delays of HDDs.

Enterprise SSDs are growing in popularity, but they are still dwarfed in comparison to HDD shipments. SSD sales are lagging for two primary reasons: cost and reliability concerns.

On a cost per gigabyte basis, SSDs still command about a 6 times price premium over HDDs. From a reliability standpoint,

(continued)

the “flash” in flash technology means that memory cells are weakened each time an electrical erasure is performed — which is required each time data is stored in flash.

However, flash vendors are tackling SSD reliability issues with sophisticated wear-leveling algorithms, and manufacturing economies of scale are driving costs down.

Today, enterprises are implementing flash and SSD technologies in hybrid storage arrays and in small-capacity (purpose-built for high performance) all-flash storage arrays.

Efficiency

The potential impact of flash technology for driving significant efficiency improvements in the data center is clear. Flash is a very flexible technology that is fabricated in a number of different form factors, such as PCI-e flash cards and SSDs. It can be deployed as a memory tier in the server, network, storage controller, or storage array. Flash can either be presented as a cache or as a standalone persistent all-flash array (see Figure 3-1).

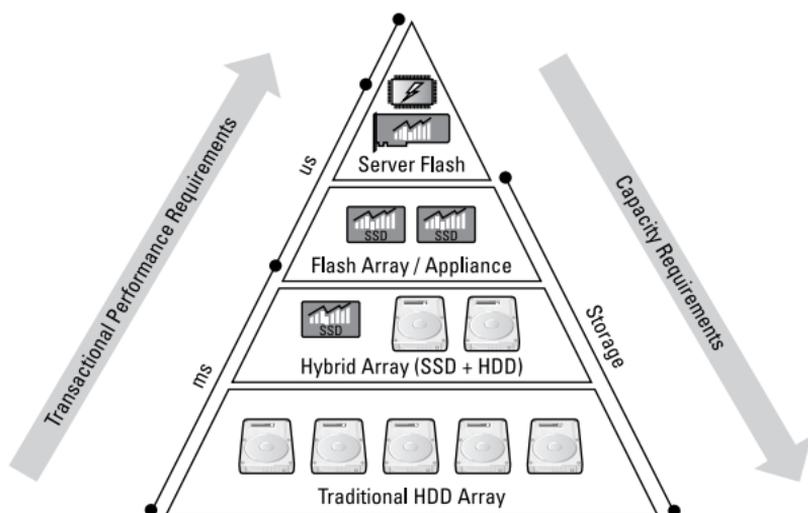


Figure 3-1: Flash technology can be flexibly deployed to address different performance and capacity requirements.

Flash storage technology helps IT organizations operate more efficiently. Rather than using techniques such as short stroking (explained in Chapter 2) for HDDs, storage administrators can use all available space with SSDs, optimize their storage, and reduce overall IT spending by utilizing storage resources more efficiently. (See the case study “Cloud services provider saves more than \$1 million with NetApp storage efficiencies.”)

Case study: Cloud services provider saves more than \$1 million with NetApp storage efficiencies

A cloud services provider needed to address scalability challenges in its storage infrastructure while maintaining acceptable service levels for more than 350,000 customers worldwide.

Challenge

Remain competitive in the hosted communication and collaboration services market with the agility and flexibility to respond to customers' needs.

Solution

Migrate customer to Microsoft Exchange Server 2010 on a seamless, integrated NetApp storage platform with performance boosted by the NetApp Virtual Storage Tier (VST).

Benefits

- ✓ Decreased total cost of ownership by 60 percent, saving more than \$1 million in IT expenditures
- ✓ Reduced storage in the Microsoft Exchange 2010 environment by 75 percent
- ✓ Reduced storage management time by 66 percent

Overview

A premier provider of hosted business communication, collaboration, compliance and security, and infrastructure solutions to more than 350,000 users worldwide must deliver its cloud services from a cost-effective infrastructure that enables high-availability, high-performance business solutions.

The storage platform supporting the company's Microsoft Exchange 2007 and Microsoft SharePoint 2007 environments had reached capacity. The platform lacked the scalability and manageability needed to upgrade these applications and migrate its customers.

The company worked closely with a technology solutions provider and NetApp Star Partner to find a robust, seamless, unified platform for managing physical and virtual resources. It supports its communication and collaboration services with NetApp FAS3270 storage systems in both its primary and secondary data centers and leverages FAS3210 and FAS3140 storage systems in its lab environment.

NetApp storage systems help the cloud services provider optimize performance and improve the customer experience with the NetApp VST driven by 512 GB NetApp Flash Cache cards (see Chapter 2) in its FAS3270 storage systems. The company leverages VST — with its intelligent caching and NetApp storage efficiency technologies — for its Microsoft Exchange and SharePoint workloads. The NetApp VST provides the ability to increase IOPS, storage efficiency, and the number of Microsoft Exchange 2010 mailboxes supported. The company achieves additional efficiencies by using NetApp deduplication, which enables it to reduce storage requirements in its Microsoft environments by 40 to 50 percent.

The NetApp storage solution helps this cloud services provider optimize its IT budget, reduce risk, and scale for growth. The midrange storage system provides a cost-effective environment for mixed workloads that protects the company's investment through its inherent scalability. The FAS3200 series storage system further enhances the high-availability NetApp architecture by enabling additional diagnostics and nondisruptive recovery.

Cost

Although the cost per gigabyte of traditional HDDs is less than the cost per gigabyte of SSDs, this is not the only measure of your storage costs.

Other important measures of cost include, for example:

- ✔ **Effective capacity:** In order to optimize IOPS, many IT storage administrators use partially filled storage arrays to reduce seek latency on HDDs. This practice also reduces the effective capacity of HDDs and significantly increases their cost per gigabyte in terms of usable storage.
- ✔ **Price per IOPS:** Inexpensive disk costs you more, in terms of productivity and lost opportunities, if it can't deliver the throughput that your most demanding applications and users require. Typical price per IOPS for HDDs is approximately \$1.25/IOPS compared to only \$0.02/IOPS for SSDs.
- ✔ **Operational costs:** Power, cooling, and data center real estate, as well as storage management and administration, are all important operational costs that must be considered.

Flash storage delivers high-performance throughput without requiring spinning platters and movable read/write heads. Instead, flash technology uses transistors to store data in an array of rows and columns (see Chapter 1).

With higher performance, your applications run faster and your users are more productive. Although some may dismiss this measure as a negligible “soft cost,” the difference between enterprise flash SSD and HDD IOPS is significant. (See the case study “NetApp FAS and Flash Cache solutions accelerate gas exploration.”)

Finally, enterprise flash SSDs can help you improve the utilization of your data center space and reduce your energy costs.

Enterprises can reduce their physical footprint in the data center by replacing high-performance, relatively lower-capacity SAS and Fibre Channel HDDs, with low-cost, high-capacity SATA HDDs and SSDs. An enterprise could easily replace 1,000 enterprise HDDs with 24 SSDs that deliver similar transactional IOPS performance (see Figure 3-2).

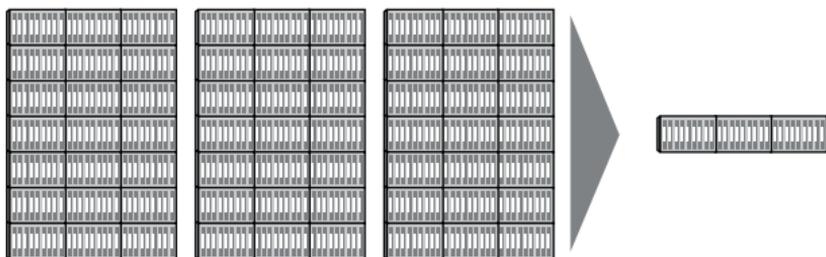


Figure 3-2: Twenty-four SSDs can provide the same IOPS performance as 1,000 enterprise HDDs.



HDDs require significant electrical power to operate their motors and spin their media. By contrast, flash storage consumes minimal power. Flash improves power consumption by 600 percent per IOPS compared to HDDs and consumes 2.5 times less power per SSD, typically 6 watts versus 16 watts for HDDs. SSDs also do not generate as much heat as HDDs, lowering your operational costs for cooling.



Enterprise flash SSDs can provide up to 60 percent savings in space and power and enhance cost per terabyte by up to 46 percent with broad enablement of SATA technology.

Case study: NetApp FAS and Flash Cache solutions accelerate gas exploration

A large energy exploration company looking to resolve significant capacity and performance challenges turned to NetApp to help deliver competitive advantage.

Challenge

Manage 200 to 300 percent annual growth and meet intensive performance requirements in a geologic modeling and seismic interpretation application environment.

Solution

Accelerate system performance and improve staff productivity with NetApp FAS6000 and FAS3000 series storage systems plus NetApp Flash Cache intelligent caching.

Benefits

- ✓ 700 percent data growth with no additional head count
- ✓ 52 percent lower cost per terabyte, 66 percent power savings, and 59 percent space savings with Flash Cache on SATA drives
- ✓ Near-instantaneous access to large data volumes

Overview

A U.S.-based independent energy company that explores for, develops, and produces natural gas, crude oil, and natural gas liquids leverages the latest technology to help maximize its

asset values of more than \$11 billion and empower approximately 4,450 employees with critical decision-making tools.

With corporate and exploration data having grown to an astounding 13 petabytes, the company must maintain an IT infrastructure that supports its changing, intense operational requirements. The company's success depends on quick access to large volumes of reliable seismic and interpretation data. In addition, the company spends millions of dollars to secure drilling rights in a particular region. To mitigate the associated financial risk, the company must give its staff ready access to highly accurate information about ideal drilling locations.

The company deployed a NetApp storage environment that scales quickly and easily and offers several important performance advantages. FAS6280 and FAS6080 storage systems at the primary and secondary data centers provide the production environment, while FAS3270 and FAS3170 storage systems support other operational areas. NetApp Flash Cache (see Chapter 2) is also an essential part of the solution, helping the company to optimize performance across the entire environment. Flash Cache speeds access to oil and gas data through intelligent caching for random read-intensive workloads.

In the shared seismic working environment, the company saw a nearly 70 percent cache hit rate and reduced latency by a factor of 10 or more with Flash Cache. This means that 70 percent of the time, data is already stored in memory, eliminating the need for the system to retrieve data from SATA drives and delivering data in seconds rather than minutes.

Running geological models that used to take 20 minutes to open and load now takes less than 5 minutes. Geoscientists can now run tests that used to take 10 minutes in just 2 minutes. And improved system performance helps the geophysics team move faster on initiatives such as developing new technologies.

Case study: NetApp EF-Series flash array accelerates Oracle database

Faced with expectations of extremely high business growth and the need to secure its own competitive advantage, an investment company set a goal of enhancing the I/O performance of its database infrastructure.

Challenge

Increase I/O performance of business-critical Oracle database.

Solution

Deploy four NetApp EF-Series flash arrays distributed over two data centers to support the investment management and business intelligence applications running on Oracle.

Benefits

- ✓ Average response times for database I/O dropped to just 0.86 millisecond, 12 times faster
- ✓ Increased performance of large-scale analytical jobs by 100 percent, now run twice as fast
- ✓ Transmitting data to neighboring satellite systems now requires only 30 percent of the time it did previously
- ✓ The time needed to recreate database cache following a restart reduced by 90 percent
- ✓ Backup times reduced by more than 75 percent

Chapter 4

Ten (Okay, Five) Strategies for Using Flash in Your Next Storage Project

In This Chapter

- ▶ Deploying flash successfully
-

Flash technology has introduced a new set of options for storage professionals. By deploying flash technology, organizations can take advantage of increased and faster throughput to significantly boost performance in mission-critical I/O-intensive applications.

For storage professionals, the challenge is this: How do you deploy flash technology to deliver the right data, at the right time, at the right cost to fit your needs? Performance, as always, is critical. And so is cost — organizations must continue to get the most bang for their buck when investing in new storage infrastructure.

With that in mind, this chapter gives you five strategies for making smart decisions about using flash for your next storage project.

Understand Your Workload

Flash is ideal for optimizing the performance of random read-intensive workloads. Although you may be targeting flash for one or more mission-critical I/O-intensive uses, such as file services, messaging, virtual infrastructure, or OLTP databases, it's critical that you also understand the data access patterns of your workloads. Here are four questions to address:

- ✔ What is the read/write mix for the workload?
- ✔ What are the typical I/O sizes?
- ✔ How much is random versus sequential?
- ✔ How large are the working sets for the application?

Employ Virtual Storage Tiering

Most automated tiering solutions attempt to continuously move data across physical storage tiers based on data access patterns and a set of predefined policies. This is a complex process that requires design trade-offs for data granularity and the timing of data movement in order to deliver an acceptable level of performance.

An alternative is to use virtual storage tiering. This self-managing approach offers real-time promotion of hot data at the smallest level of granularity, without the need for complex data classification and movement. With virtual tiering, you simply choose the default

media tier you want for a volume or LUN, such as SATA, Fibre Channel, or SAS, and data from the volume or LUN is automatically promoted to flash-based media on demand.

Evaluate Price/Performance at a System Level

When evaluating flash technology, it can be tempting to make comparisons at a component level and then extrapolate the results. Don't fall into that trap. For example, if your system is optimized for write performance, flash technology may not be the most appropriate storage medium. Remember, the primary benefit of flash is typically the acceleration of read performance. What matters most at the application layer is not which subsystem handles the reads and writes but at what speed and cost the data is delivered.

Leverage All Flash Where It Counts

All-flash arrays have a much higher cost per gigabyte than hybrid or all-HDD systems, so it's important to understand when to leverage these systems for maximum impact. All-flash arrays are best suited for latency-sensitive applications where improved responsiveness results in improved productivity, faster time to market, better decision making, or higher revenues. Examples include transactional workloads (OLTP), real-time analytics (OLAP), decision-making systems, and technical computing. All flash is also great for improving end-user experience in virtual desktop infrastructures with very high IO-per-desktop requirements.

Deploy Flash with Capacity-Optimized Drives

You should consider NetApp Flash Cache (discussed in Chapter 2) and Flash Pool in situations where a disk array provides sufficient capacity but additional I/O is needed. Testing shows that Flash Pool enables a system with SATA drives to match the performance of a system with SAS drives, while also providing 50 percent more storage capacity, 46 percent lower cost per terabyte, and a reduction in electrical power requirements of approximately 27 percent.

Deploy flash storage for mission-critical applications

Flash technology brings unprecedented speed, reliability, and efficiency to enterprise data centers. This book shows you how and where to deploy flash in your storage infrastructure.

- *Simplify and optimize your storage with virtual storage tiering — automatically copy “hot” data to cache when needed*
- *Use flash to improve storage performance — perform read and write operations in microseconds and increase IOPS to tens of thousands, even millions*
- *Reduce operational costs — including power, cooling, and data center real estate*

Lawrence C. Miller, CISSP, has worked in information security and technology management for more than 15 years. He has written more than 50 *For Dummies* books on numerous topics, including *CISSP For Dummies*.



Open the book and find:

- Which flash technologies are best for your storage challenges
- How to replace SAS hard drives with SATA and flash
- Why solid state drives are faster and more reliable than disk
- How flash can improve user productivity

Making Everything Easier!™

Go to **Dummies.com®** for videos, step-by-step examples, how-to articles, or to shop!

FOR
DUMMIES[®]
A Wiley Brand

ISBN 978-1-118-99994-3
Book not for resale