

Compliments of  NetApp®

NetApp Special Edition

All-Flash Storage

FOR DUMMIES®

A Wiley Brand

Learn to:

- Optimize storage performance
- Match your workloads with the best all-flash solution
- Reduce the footprint of your storage infrastructure

Lawrence C. Miller, CISSP



Leading organizations worldwide count on NetApp for software, systems, and services to manage and store their data. Learn more about the NetApp portfolio of all-flash solutions and its vision for the future of data management at netapp.com.

***All-Flash
Storage***
FOR
DUMMIES[®]
A Wiley Brand

NetApp Special Edition

***All-Flash
Storage***
FOR
DUMMIES[®]
A Wiley Brand

NetApp Special Edition

by Lawrence C. Miller, CISSP

FOR
DUMMIES[®]
A Wiley Brand

All-Flash Storage For Dummies®, NetApp Special Edition

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

Copyright © 2016 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, For Dummies, the Dummies Man logo, 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, ONTAP, Snapshot, SolidFire and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netappmlist.aspx>. 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 978-1-119-32893-3 (pbk); ISBN 978-1-119-32897-1 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgments

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

Project Editor: Jennifer Bingham

Acquisitions Editor: Katie Mohr

Editorial Manager: Rev Mengle

Business Development Representative:
Karen Hattan

Special Help from NetApp: Larry Freeman,
Rob McDonald, Anjali Acharya, Jim Lyons

Production Editor: Antony Sami

Table of Contents

Introduction	1
Foolish Assumptions	1
About This Book	2
Icons Used in This Book.....	2
Where to Go from Here	3
Chapter 1: All-Flash Storage Technology — What and Why	5
Flash Concepts	5
SLC, MLC, and Lots of TLC.....	6
The Evolution of Flash Technology.....	8
Read-only memory	8
Erasable programmable read-only memory	9
Compact Flash	10
Flash in the Enterprise	12
Chapter 2: All-Flash Storage Technology — Where and How	15
Different Platforms for Different Needs	16
All-flash arrays for 2nd Platform environments	16
All-flash arrays for 3rd Platform environments	17
All-flash arrays for applications that need extreme performance.....	17

Flash-to-Disk-to-Cloud	19
As-a-Service Environments	20
Scale-out	21
Guaranteed performance and QoS.....	21
Automated management	22
Extreme Performance.....	22

Chapter 3: Recognizing the Benefits of All-Flash Storage Systems 25

Performance	25
Flexibility.....	27
Simplicity	29
Total Cost of Ownership	32

Chapter 4: Ten (Okay, Five) Strategies for Using Flash in Your Next Storage Project..... 37

Clarify Performance Claims	38
Ask for a Test Drive	39
Get It in Writing	39
Focus on TCO Instead of the Purchase Price.....	39
Use an Extended Warranty for Investment Protection	40

Introduction



Flash technology has made quite a splash in the storage industry. Offering superior speed and reliability when compared to disk-based and hybrid flash systems, all-flash storage has proven to be cost-effective for an increasingly wide range of enterprise applications.

All-Flash Storage For Dummies, NetApp Special Edition, explores the expanding role of all-flash storage in the enterprise. Although all-flash systems were originally used to accelerate performance for dedicated, mission-critical applications, they're now replacing disk systems as the preferred storage infrastructure for virtualized environments. Meanwhile, new scale-out solutions have emerged that streamline IT operations and reduce the cost of data centers designed to run the cloud services of the future.

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: All-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 a “flash-first” approach your storage infrastructure.

Chapter 2: All-Flash Storage Technology — Where and How: Here, you find out about NetApp's flash storage solutions: All-Flash FAS (AFF), SolidFire, and the EF-Series.

Chapter 3: Recognizing the Benefits of Flash Storage Systems: This chapter tells you how all-flash storage can benefit your organization, including some areas you may not have considered.

Chapter 4: Ten (Okay, Five) Strategies for Using All-Flash Systems in Your Next Storage Project: Finally, in that classic *For Dummies* style, I give you some useful tips on how to deploy 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 *All-Flash Storage For Dummies*, NetApp 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

4

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

All-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
- ▶ The flash technology “tipping point”

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 with 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 Lots of TLC

Flash technology has many advantages, but in the early days of flash, vendors had to overcome a couple of inherent weaknesses before the technology could really take off:

- ✔ 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 wasn't very fast. In fact, early flash-based solid state drives (SSDs) were slower than high-performance HDDs under certain conditions, such as large sequential write operations.



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 rewritten. 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, new technologies have been developed and they're already leading some companies to replace HDDs altogether.

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

To get the durability of SLC flash in MLC flash without the extra cost, flash-component manufacturers began to offer MLC flash with sophisticated wear-leveling and bad-block management algorithms. These innovations enable lower cost MLC and TLC flash to be used in enterprise SSDs without sacrificing reliability.



More recently, a new technology has been developed that provides dramatically higher flash capacities, but without the trade-offs associated with shrinking the cell geometry. Instead, flash capacity is increased by stacking the cells in multiple layers using a process called 3D-NAND (sometimes referred to as

V-NAND). By maintaining the size of the individual cells, the performance and endurance issues of cell-to-cell interference can be avoided. Because of these advantages, flash-component suppliers have begun ramping up their production of vertically stacked MLC and TLC flash components. These components are increasing the densities of SSD drives while also lowering their cost per bit.

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 capability to erase and rerecord single memory cells, unlike earlier EEPROMs that had to be completely erased before any new data could be stored.

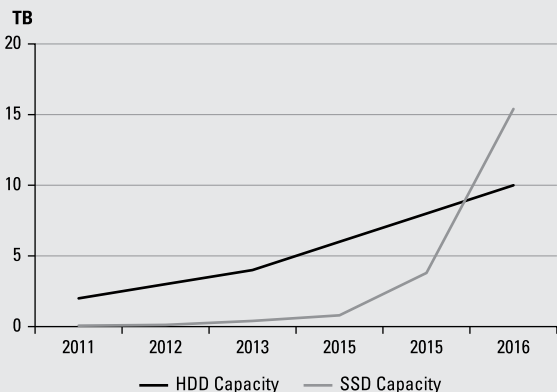
Thus a new era of storage was born and the rest, as they say, is history! As the demand for consumer electronics took off, the cost of flash technology plummeted to the point where storage vendors took notice, and many IT professionals are now asking “Will flash completely replace traditional hard drives?”

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, global unit sales peaked in the 2010 time frame and are now beginning to decline at an accelerating rate. Precision motors and actuators are subject to mechanical wear and don't perform indefinitely. Spinning drive platters 24 hours a day but accessing data only occasionally is an inefficient use of electrical energy. Based on current trends, the capacity, reliability, and economic ceilings for HDDs will soon be reached (see the accompanying figure).

The reason the HDD has survived so long is that nothing better has come along worthy of replacing it — until now.



Source: NetApp

Flash in the Enterprise

Flash was originally introduced as a premium storage performance tier. By substituting solid-state performance for mechanical disk drives, flash users could realize 10 to 20 times performance improvements for critical databases and virtual desktop environments, and they could improve performance for the majority of read operations in hybrid designs that combined SSDs and HDDs. But the adoption of all-flash systems was limited because SSD capacity was simply more expensive than HDD capacity.

Recently, however, the price of flash media has dropped dramatically. Semiconductor costs are largely driven by volume, and flash technology has supplanted hard disks for popular consumer devices such as mobile phones, tablets, and notebooks. In other words, enterprise storage customers are benefitting from the growing use of flash in consumer technologies.

At the same time, the process technology improvements in flash component design have enabled more information to be crammed into each device, leading to improved SSD density. SSD capacities have now surpassed those of the largest hard disk drive, and the capacity advantage is expected to keep increasing for the foreseeable future. These high-capacity SSD devices provide much more efficient use of power, cooling, and rack space compared with HDDs (see Figure 1-1).

One petabyte of storage capacity
with 10K high-performance HDDs

vs. One petabyte of storage capacity with
15TB SSDs and 4:1 storage efficiency




Figure 1-1: Rack space for 1PB of effective capacity.

Source: NetApp

SSD devices offer the following benefits:

- ✓ Less than one millisecond of latency performance for all data
- ✓ A much smaller footprint than an HDD
- ✓ Much less power consumption
- ✓ No tuning for performance
- ✓ Significant savings in support costs

New compression and deduplication technologies further reduce flash costs by cutting the absolute capacity requirements. Although disk drive technologies relied on “post-process” storage efficiencies, all-flash systems can perform compression and deduplication operations in real time. These real-time storage efficiencies can reduce the amount of flash capacity required by a factor of four times or more, depending on the application.

In his best-selling book, Malcolm Gladwell described the rapid adoption of powerful new ideas, products, and norms when conditions are right and labeled it *The Tipping Point*. A tipping point requires an environment ready for change. Today's enterprise IT environments, which are characterized by a scarcity of resources and economic pressure to control costs, are ready for change at precisely the same time that flash technology is making rapid improvements. As a result, all-flash arrays have reached a tipping point and are now replacing both disk and hybrid flash architectures for an increasing range of primary IT workloads.

Chapter 2

All-Flash Storage Technology — Where and How

In This Chapter

- ▶ Choosing the right all-flash platform
 - ▶ Eliminating silos of all-flash storage
 - ▶ Building scale-out platforms (like Google!)
 - ▶ Architecting solutions for extreme performance
-

All-flash storage systems can improve both I/O performance and efficiency for the majority of your primary workloads, including databases, virtual servers, and next-generation cloud services. Chapter 1 explains what flash technology is and why it's rapidly replacing disk systems in the enterprise. In this chapter, you get acquainted with NetApp's portfolio of all-flash storage solutions and learn how each is designed to meet the needs of a specific type of application environment.

Different Platforms for Different Needs

When selecting all-flash storage systems, it helps to start by considering your application environment, and these tend to fall into one of three categories:

- ✓ **2nd Platform environments** that include a mix of traditional SAN and NAS workloads that share a common infrastructure.
- ✓ **3rd Platform environments** that include new next-generation applications designed for a cloud infrastructure.
- ✓ **Dedicated applications** that require extremely low latency or extremely high throughput.

All-flash arrays for 2nd Platform environments

The analyst firm IDC coined the terms *2nd Platform* and *3rd Platform* to describe the shift from traditional client/server computing to a new mix of applications optimized for social, mobile, analytics, and cloud (sometimes referred to as *SMAC*).

For many companies, and especially those that have been around for a while, the majority of their applications still run on-premises and fall into the 2nd Platform category. These applications typically run on virtual servers and share a common hardware infrastructure in order to minimize IT costs.

Most 2nd Platform environments require all-flash arrays that can deliver robust data management in addition to sustained high performance, including capabilities such as:

- ✓ Inline storage efficiencies
- ✓ Replication to non-flash platforms, such as disk systems and public clouds
- ✓ Integration with enterprise software from Microsoft, Oracle, SAP, and VMware
- ✓ Interoperability with a mix of legacy server and storage hardware

All-flash arrays for 3rd Platform environments

Many 3rd Platform environments support applications that are designed to run in the public cloud or on top of a cloudlike infrastructure where software is delivered as-a-service, either to internal users or to external customers.

These environments require scale, agility, and economics similar to public cloud infrastructures run by Amazon or Google. Although 3rd Platform environments require many of the same storage capabilities as 2nd Platform environments, they also require additional features, including:

- ✓ Seamless scale-in and scale-out
- ✓ An ability to quickly scale capacity and performance in small or large amounts
- ✓ Sophisticated quality of service (QoS) controls
- ✓ Software-defined architectures with advanced automation capabilities

All-flash arrays for applications that need extreme performance

For some applications, such as mission-critical online transaction processing, databases, and data analytics,

what matters the most is extreme performance. Anything that slows down the I/O path to the storage media is considered unacceptable. Even the small amount of overhead introduced by inline storage efficiencies can introduce delays and cause problems for these applications.

These applications tend to require dedicated hardware in order to ensure the highest levels of performance and often run on bare metal servers rather than virtual servers due to their requirements for speed. Although many all-flash storage systems can support latencies as low as one millisecond, applications that require extreme performance require response times that are even faster and are measured in microseconds rather than milliseconds.



When selecting an all-flash array, start by considering your application requirements. Do you need support for a 2nd Platform environment, a 3rd Platform environment, or a dedicated application that requires blazing fast performance above all else? See Table 2-1.

Table 2-1 Criteria for Evaluating All-Flash Arrays

<i>Application Environment</i>	<i>Typical Use Cases</i>	<i>Key Buying Criteria</i>
2 nd Platform	Virtualized applications	Enterprise data management
3 rd Platform	Next-generation cloud services	Scale-out, service oriented architecture
Dedicated applications	Online transaction processing	Extreme performance with low latency

Flash-to-Disk-to-Cloud

The first generation of all-flash arrays were almost exclusively designed to accelerate dedicated applications, such as databases and virtual desktops. These systems typically supported the Fibre Channel SAN protocol and were operated as storage silos without an ability to support NAS-based filesystems or replicate data to alternate storage tiers, such as HDD-based systems, hybrid arrays, or public clouds.

These early designs worked well for the target applications, but they lacked basic enterprise features such as active-active controller pairs, robust storage-efficiency technology, and integrated data protection. They were also inefficient to operate for several reasons:

- ✔ Performance and capacity were stranded and could not be shared across workloads.
- ✔ The arrays needed to be managed separately.
- ✔ Protection copies could not be stored on lower-performance, lower cost storage tiers.

NetApp All-Flash FAS (AFF) systems provide a high-performance solution for consolidating multiple SAN and NAS workloads on a single, unified architecture. With AFF systems, your investment is protected if your performance and capacity needs change or if your IT strategy evolves to encompass data management across a hybrid cloud.

Here are some of the benefits of NetApp AFF systems:

- ✔ They eliminate performance silos and seamlessly integrate with hybrid flash systems in an ONTAP environment, enabling multiple performance tiers.

- ✔ Workloads can be transparently moved to the storage tier that best meets your applications' requirements.
- ✔ You can consolidate a wide range of workloads, with support for millions of IOPS and hundreds of petabytes of effective capacity in a single cluster.
- ✔ They're Data Fabric ready, so you can move data to and from the cloud to optimize performance and your return on investment.
- ✔ Interoperability with existing infrastructure enables support for a broad set of applications and provides the flexibility to add or change workloads as needed.



Data Fabric is the NetApp vision for a storage architecture that enables data to move seamlessly between on-premises infrastructure and multiple public cloud services within a common data management framework.

As-a-Service Environments

Cloud architectures extend the virtues of virtualization by offering resources on-demand and as a service. Cloud computing allows an enterprise to maximize network access to a shared pool of elastic, scalable resources.

NetApp SolidFire all-flash systems are designed to deliver the scale required for cloud-based services, guarantee storage performance, and enable complete system automation.

Scale-out

In a scale-out design, pools of storage (GB, IOPS) are shared resources. The capability to cluster these resources provides incremental, on-demand scaling. IT teams can purchase and manage only what the business currently needs — even one node at a time.

SolidFire scale-out storage provides the flexibility to independently and nondisruptively scale both capacity and performance in a predictable linear pattern. This means that businesses can distribute data and traffic across all of the nodes as the scope of their data services changes over time. IT teams gain efficiencies by leveraging cloudlike elasticity within the storage layer and integrating tightly with other services, resources, and management platforms.

Guaranteed performance and QoS

An ability to guarantee firm performance service levels across a storage infrastructure is a key differentiator for cloud service providers and enterprise clouds alike.

Guaranteed performance means that QoS controls are enforced at a granular level by allocating a specific amount of storage resources to each application. Every application provisioned is assigned a guaranteed range of IOPS and those allocated IOPS are respected consistently, regardless of any other application activity, capacity level, or I/O pattern. This consistency creates predictable performance based on the requirements of the business, while isolating one workload from another and ensuring performance in the event of hardware or software faults.

SolidFire enables you to define the following QoS policies:

- ✔ **Min IOPS:** I/O operations per second that are always available to a volume. This ensures that performance guarantees are maintained regardless of system condition or application activity.
- ✔ **Max IOPS:** I/O operations per second that a volume can process over a sustained period of time.
- ✔ **Burst IOPS:** I/O operations per second that a volume will be allowed to process during a spike in demand. This is particularly effective for uneven and latency-sensitive workloads.

Automated management

In cloud-scale storage deployments, automation is the key to efficiency. SolidFire all-flash storage enables a transition from a capacity-provisioned paradigm to an automated, programmable, and policy-driven environment that enables a next-generation data center design. It eliminates device-centric and endpoint-centric administration and fully enables multitenant data management and protection. Web-scale principles of automation, QoS, and API-based access help orchestrate workloads and their integration with service-driven delivery solutions. These attributes enable IT teams to deliver higher business value for the enterprise.

Extreme Performance

The EF-Series all-flash arrays are designed specifically for high-speed transactional applications that demand high IOPS and consistent low latency. For example, a

single EF system can deliver over 825,000 sustained IOPS with response times measured in microseconds. The extreme performance provided by EF-Series arrays can:

- ✔ Deliver faster, actionable results from OLTP and OLAP environments
- ✔ Significantly improve the performance of data analytics applications
- ✔ Improve the user experience for customer-facing and decision-support systems

For comparison, a single EF-Series all-flash system in a 2U enclosure can deliver the performance of over 1,000 15K RPM disk drives while requiring just a fraction of the rack space, power, and cooling! With a significant reduction in space and power consumption, all-flash systems that provide extreme performance can significantly improve the overall efficiency of IT operations while delivering the super-fast response times required by mission-critical applications. Figure 2-1 shows the entire NetApp portfolio of all-flash solutions.




AFF	SolidFire	EF
General purpose virtual infrastructure	Next generation cloud infrastructure	Optimized for low-latency workloads
		
Data Services	Scale-out	Extreme Performance

Figure 2-1: NetApp all-flash solutions.

Source: NetApp

Chapter 3

Recognizing the Benefits of All-Flash Storage Systems

In This Chapter

- ▶ Improving storage performance with flash
 - ▶ Increasing your infrastructure flexibility
 - ▶ Simplifying IT operations with all-flash systems
 - ▶ Measuring the total cost of your storage systems
-

In this chapter, you find out how IT teams can benefit from all-flash storage systems, as well as how IT organizations in different industries have successfully deployed all-flash storage and are now enjoying its benefits.

Performance

Flash technology improves overall application performance. Enterprise all-flash storage systems perform read and write operations in microseconds and provide IOPS performance that ranges from the hundreds of

thousands into the millions. By comparison, HDD-based systems are an order of magnitude slower, with response times measured in milliseconds. Similarly, the read throughput for SSDs now surpasses the read throughput for HDDs by a wide margin (see Figure 3-1).

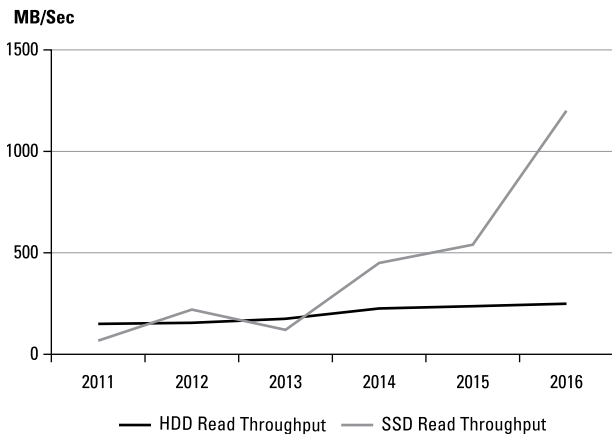


Figure 3-1: HDD and SSD read throughput trends.

Source: NetApp

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.



All-flash storage can improve response times and increase I/O throughput up to an order of magnitude (10x) compared with HDD-based systems.

See the case study at the end of the chapter, “Gaining Insights from a Mountain of Data with NetApp EF-Series” to get an idea of the business impact companies are experiencing as they transition from disk systems to all-flash.

Flexibility

Flexibility is an attribute of all-flash storage systems optimized for the 2nd Platform environments covered in Chapter 2. In fact, the robust data services these systems provide are required in order to make a transition from traditional disk-based storage to an all-flash storage infrastructure. In order to maximize operational benefits, all-flash arrays deployed into 2nd Platform environments need to integrate with existing infrastructure rather than create new data silos that increase management complexity.

This integration requires all-flash systems that can support both SAN and NAS workloads, provide an ability to scale-out to support all of your applications on a common all-flash foundation, and help you seamlessly move data to and from externally hosted cloud systems, if needed.

See the case study “University moves all of its applications to all-flash storage” for one such example.

Case study: University moves all of its applications to all-flash storage

By moving all of its on-premises applications to All Flash FAS, a U.S. university secured enough performance headroom to last for the next five years, regardless of fluctuating and unpredictable IT budgets.

Challenge

To stay competitive, the university needs to support new technology initiatives, including an increase in video content and online classes, but steadily increasing workloads were already impacting performance and causing response time delays for students, faculty, and IT staff.

Solution

The university migrated its entire VMware virtual server environment to all-flash storage.

Benefits

- ✔ Improves productivity and learning by accelerating application performance by a factor of three
- ✔ Supports distance-learning initiatives with high-performance online courses
- ✔ Enhances campus safety by supporting new IP video cameras in addition to I/O-intensive applications

Overview

To provide consistent high performance and avoid downtime for students and faculty, the university migrated its entire VMware virtual server environment to NetApp

All Flash FAS systems. The university can maintain 24/7 operations even in the event of a hurricane or other disaster by using all-flash storage to accelerate enterprise applications, while using cost effective spinning disks to store offsite disaster recovery copies.

By removing bottlenecks to performance for its core educational applications, the university can more effectively compete for top students and faculty. The school can also support and grow its distance-learning initiatives by offering interactive, data-rich online courses. And it can invest IT staff time in delivering creative solutions instead of the tedious and time-consuming task of managing performance.

“Most universities are going to find themselves moving to all-flash systems within the next few years because application performance requirements are increasing and user expectations are rising. I’m glad we made the leap now.”
— University CIO

“We’ve seen at least a threefold performance improvement for all of our applications by moving our VMware environment and Oracle databases to NetApp All Flash FAS. Everybody has a smile on their face.”
— Director of University Computing Services

Simplicity

The ability of flash technology to drive significant operational improvements in the data center is clear. When storage performance no longer needs to be treated as a scarce resource that must be closely managed, IT personnel are freed up across the entire organization.

Database, application, and infrastructure administrators can all reclaim the time spent planning for and managing storage performance.

However, there are additional opportunities for all-flash systems to simplify IT operations. Out-of-the-box templates can streamline the configuration of new systems so that IT generalists can set-up and provision storage without the need for specialists. QoS policies can be applied to guarantee performance for applications and users, eliminating yet another management headache. Perhaps even more important for many organizations, lifecycle operations such as adding and removing storage hardware and load balancing performance, can all be performed without disrupting critical business applications.

See the case study “Delivering a first-class cloud experience with all-flash storage” for a real world example.

Case study: Delivering a first-class cloud experience with all-flash storage

A multinational internet service provider needed to deploy a new cloud offering that could deliver enterprise-grade storage services and improve ease-of-use.

Challenge

The cloud service was originally built using a traditional storage architecture that offered multiple performance tiers, each of which was based on different architectures and offered at different price points. This created extra complexity, both for end customers and for the operations team that managed the internal infrastructure.

Solution

After assessing a number of vendors, the service provider chose SolidFire for its next generation storage platform.

Benefits

- ✓ Reduced storage infrastructure complexity
- ✓ Ability to easily scale-out storage performance and capacity
- ✓ Guaranteed performance service levels for all of the cloud provider's customers

Overview

Storage is a critical part of a cloud service infrastructure. This service provider was looking for a trusted storage partner and after a thorough evaluation, realized the SolidFire platform would greatly simplify its end-users' experiences while simultaneously reducing internal complexity.

"Always guaranteeing performance in spite of 'noisy neighbors' presented us with a constant challenge. While it was easy enough for us to add storage nodes to our original SAN setup, adding new customers made the problem of juggling resource demands incredibly complex."
— Head of Product Management for Cloud Service Provider

With the SolidFire solution, each time a storage node is added to the infrastructure, the cloud service benefits from predictable performance and capacity gains. SolidFire's QoS guarantee has also allowed the service provider to safeguard against "noisy neighbor" problems

(continued)

(continued)

where a user could experience a performance lag because one application affects another's performance.

Instead, each application is provided with the resources it requires without impacting the performance of the others. As a result, the service provider's operations team no longer needs to constantly manage the storage infrastructure when end customers spin-up particularly resource-intensive applications.

"Our aim is to make the cloud easy — that's our promise to our customers. We take the complexity of building a cloud away from end-users, and focus on giving them a first-class cloud experience — low latency, maximum performance, and a simple pricing structure."

— Head of Product Management for Cloud Service Provider

Total Cost of Ownership

Although the per-gigabyte cost of physical capacity for HDD-based systems may be less than the per-gigabyte cost for all-flash solutions, this turns out to be a flawed approach for comparing the two alternatives.

In order to compare the total cost of ownership (TCO), there are several other items that should be considered including:

- **Effective capacity:** This is the amount of logical capacity available to store application data after storage efficiencies are applied. At a minimum, your calculations should include the inline

efficiencies available with an all-flash system, such as data compression and deduplication. In many cases, it is also appropriate to include Snapshot copies and clones if you expect to use them as part of your normal operations.

- ✔ **Application integration:** The cost of application integration is often overlooked when evaluating primary storage. However, the costs and capabilities for provisioning, protecting, and recovering application-consistent copies of data can vary widely across all-flash offerings.
- ✔ **Operational costs:** Power, cooling, and data center real estate are all important operational costs that must be considered. According to user surveys, all-flash storage systems can provide huge savings in these areas when compared with disk-based systems.
- ✔ **Software licenses:** The consolidation of workloads onto an all-flash platform can also generate software savings. For example, database software licenses and maintenance fees can be reduced when all-flash systems enable fewer servers and CPU cores to deliver the desired level of application performance.

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 as a “soft” benefit, the cumulative productivity gains from decreasing wait times for

hundreds (or many thousands) of employees and/or customers can be material.

Finally, enterprise flash SSDs can help you improve the utilization of data center space and reduce your energy costs. For data centers where rack space is either very expensive, or at such a premium that the only alternative is to expand into a new facility, the space savings alone can be enough to justify an investment in all-flash storage.



All-flash storage systems provide big savings in power, cooling, and rack space along with savings in time and money for IT administration when compared with disk-based storage systems.

Case study: Gaining insights from a mountain of data with NetApp EF-Series

Double-digit business growth, coupled with its customers' increasing appetite for big-data analytics put heavy demands on the IT infrastructure for a provider of real estate property information.

Challenge

Speed delivery of data to help customers make better-informed property and financial decisions.

Solution

Strategically deploy an all-flash storage solution to meet the demands of high-performance databases and analytics.

Benefits

- ✓ Reduced database transfer times by over 70%
- ✓ Data transfer rates in excess of 1.2M records per second
- ✓ Enabled 99.9999% (“six nines”) availability

Overview

The company continuously collects and manages data and imagery, including maps and high-resolution images, from more than 130 feeds. Each day, its systems ingest large volumes of data to feed a database of more than 500 million property decision points. Success depends on how quickly the company can provide quality reports to its customers, which include real estate and financial services companies as well as homeowners, investors, and government organizations.

“Our customers care about three things: speed, quality, and accuracy.”

— Head of Information and Communication Technologies

“All of our products, services, and systems are there to give customers access to the information and analytics they need to make confident decisions about the transactions they’re performing.”

— General Manager of Technology

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 for the majority of their business applications.

For storage professionals, the challenge is this: When all-flash performance and reliability are “table stakes” and available from nearly every vendor, how do you select the solution that can best meet your needs? Real-world performance, as always, is still important. And so is cost — organizations must continue to get the most bang for their buck when investing in new storage infrastructure.

However, there are a number of other things that should be factored into your selection process. This chapter gives you five strategies for making smart decisions about all-flash systems for your next storage project.

Clarify Performance Claims

Flash is ideal for optimizing performance, but there can be wide performance variations across vendor systems, and even wide variability within a single system depending on a number of factors. In general, you should seek solutions that can demonstrate consistent low latency based on third-party benchmarks that simulate real-world workloads, such as SPC-1 and TPC-E. You should also take claims about “maximum IOPS” with a grain of salt unless your vendor can provide detailed answers to the following questions:

- ✓ What was the I/O size used to run the test?
- ✓ How does this compare with the I/O size for my applications?
- ✓ What was the read/write mix?
- ✓ Were the I/Os random or sequential?
- ✓ And most important of all: What was the latency at increasing IOPs intervals, and at what IOPs level did the latency exceed one millisecond?

Without the answers to these questions, you will be unable to make a valid apples-to-apples comparison across competing solutions.

Ask for a Test Drive

Although benchmark data is useful for making comparisons, the real litmus test is to run an all-flash solution using your own applications and data. Vendors typically offer a limited time test drive program and some will even let you use their own facilities to run sophisticated proof-of-concept (POC) tests.

However, beware of vendors that use sales tactics to try and avoid a hands-on evaluation, or who insist on using nonstandard configuration settings or their own special data sets to demonstrate performance or storage efficiency savings. These are all red flags that a vendor's performance claims may not be achievable in a real-world environment.

Get It in Writing

When evaluating all-flash systems, you're likely to hear performance and efficiency claims that are simply too good to be true. If you're unable to verify (or disprove) such claims by testing with your own applications and data, then you should request that your vendor back it up in writing, including the remediation that will be provided to your company for failure to deliver as promised. Reputable vendors will back-up their claims in writing.

Focus on TCO Instead of the Purchase Price

In Chapter 3, you looked at the cost savings that can be achieved with all-flash arrays over and above the cost of physical storage capacity. When comparing all-flash

systems to HDD-based or hybrid flash systems, you should always take the time to identify the potential savings in these other areas and calculate the total cost of ownership for the period of time you plan to operate the systems.

Use an Extended Warranty for Investment Protection

If this is your first experience with an all-flash system, you may be concerned about flash wear out, a topic we touched on in Chapter 1. However, you will discover that most vendors are so confident in the long-term reliability of their all-flash systems that they're happy to provide you with an extended warranty. Other options, such as free controller upgrades, may also be available if you're willing to commit to a long-term service contract. Remember, you'll need to factor all of the costs into your total cost of ownership calculations in order to make an apples-to-apples comparison across vendors. For example, you may discover that the potential savings from some of the "free" controller upgrade options are more than offset by higher annual service charges.

Deploy all-flash storage for all of your applications

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

- *Simplify storage management — replace disk systems with all-flash*
- *Use flash to improve storage performance — perform read and write operations in microseconds and increase IOPS to hundreds of thousands, even millions*
- *Reduce operational costs — including power, cooling, and data center real estate*



Open the book and find:

- Which flash solutions are best suited for your storage challenges
- How you can use flash for high-capacity storage requirements
- How flash can simplify IT operations

Go to Dummies.com[®] for more!

FOR
DUMMIES[®]
A Wiley Brand

ISBN: 978-1-119-32893-3
Not for resale

WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.