

Making Everything Easier![™]

Pure Storage Edition

All-Flash Data Centers

FOR
DUMMIES[®]
A Wiley Brand

Brought to you by



Scott D. Lowe



About Pure Storage

Pure Storage helps companies push the boundaries of what's possible. The company's all-flash based technology, combined with its customer-friendly business model, drives business and IT transformation with Smart Storage that's effortless, efficient, and evergreen. Pure Storage offers two flagship products: FlashArray//m, optimized for structured workloads, and FlashBlade, ideal for unstructured data. With Pure's industry leading Satmetrix-certified NPS score of 83, Pure customers are some of the happiest in the world and include organizations of all sizes, across an ever-expanding range of industries.

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

Pure Storage Edition

by Scott D. Lowe

FOR
DUMMIES[®]
A Wiley Brand

All-Flash Data Centers For Dummies®, Pure Storage Edition

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

Copyright © 2017 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. Pure Storage, Pural, the P Logo, Evergreen, Elasticity, FlashBlade, FlashStack, Hello Possible, and Accelerate Possible are trademarks or registered trademarks of Pure Storage, Inc. in the U.S. and 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 978-1-119-32923-7 (pbk); ISBN 978-1-119-32924-4 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgments

We're proud of this book and of the people who worked on it. For details on how to create a custom *For Dummies* book for your business or organization, contact info@dummies.biz or visit www.wiley.com/go/custompub. For details on licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

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

Project Editor: Carrie A. Burchfield

Acquisitions Editor: Katie Mohr

Editorial Manager: Rev Mengle

Business Development Representative:

Karen Hattan

Production Editor: Siddique Shaik

Table of Contents

Introduction	1
About This Book	1
Foolish Assumptions	1
Icons Used in This Book.....	2
Beyond the Book.....	2
Chapter 1: Containing Storage Costs	3
Calculating the TCO of Traditional Storage.....	3
Discovering the All-Flash Ingredients that Reduce TCO.....	5
How data reduction amplifies available capacity.....	6
Facilities improvements: power and cooling.....	7
Data center storage footprint and cabling	7
Eliminating Storage Software Licensing	
Headaches and Costs	8
The Importance of Smart Storage.....	9
Chapter 2: Reaching Simplicity and Manageability	11
Why Simplicity is The New Black: Scaling the Business	11
How Does Storage Simplicity Manifest in an	
Organization?.....	12
Reducing procurement complexity.....	13
Repurposing IT resources	13
Reducing the data center maintenance burden	13
Tuning Flash Storage for Application Workloads	14
Busting Flash Reliability Myths.....	15
Chapter 3: Supporting Virtualization and Consolidation	17
Identifying Key Storage Performance Challenges	18
Performance concerns.....	18
Lack of budget.....	19
Purchasing additional infrastructure.....	19
Complexity.....	20
Solving Virtualization and Consolidation Challenges	20
Reduced latency	20
Liberating server resources to improve	
VM density	21
Deduplication	21
Management integration with the hypervisor	22

Chapter 4: Predicting and Modeling Future Storage Consumption and Costs	23
Understand Long-Term Storage Economics	23
Discover Inline Deduplication's Cost Benefits	25
Guaranteed data reduction	26
Predictable performance	27
Chapter 5: Achieving Scalability and Agility	29
Using Simplicity and Operational Effectiveness for Scalability and Agility	29
Consolidate storage-related services	30
Adopt converged infrastructure	30
Manage with the right tools	31
Stop worrying about upgrades	32
Driving Scale with Predictive Application Performance.....	33
Leveraging APIs for Automation	33
Supporting Data Growth	34
Chapter 6: Addressing the Cloud	35
Looking at Flash and Cloud as Key Data Center Disruptors	35
Understanding Cloud Characteristics versus Destination.....	36
Looking to All-Flash Data Centers as the Foundation for Enterprise Cloud Strategies	37
Simple scaling.....	38
Predictable performance	38
Always-on.....	38
Easy management from anywhere	39
OpEx cost models	39
Integration capabilities	40
Chapter 7: Ten Reasons You Need All-Flash Storage	41
Flash Is Fast . . . Really Fast	41
Flash Reduces Power and Cooling.....	42
Flash Increases Usable Capacity	42
Flash Provides a Comprehensive API.....	42
You Can Converge All the Things	43
All-Flash Architecture Effectively Ends the Forklift Upgrade.....	43
All-Flash Reduces Your Storage Footprint.....	44

Introduction

Storage. Many organizations lament the fact that traditional storage is plagued with complexity and waste, and they've looked to cloud IT for answers. Those that choose to adopt smart storage solutions discover the answers they've been looking for. Today, all-flash systems take a strategic role in the data center and provide companies with cloud-like productivity, efficiency, and economics.

The era of smart storage systems is upon us and therein lies an opportunity to transform the entire data center and, believe it or not, businesses as a whole. Today, with the right all-flash storage system, it's easier than ever to refocus IT on the needs of the business and achieve the agility and efficiencies promised by cloud IT.

About This Book

In this astounding 48-page book, you go on a journey that provides a look at some of the meta trends shaping the emerging IT department of the future. Here's a sneak peek: Storage is undergoing a fundamental transformation. Storage administrators and CIOs discover new workload and efficiency capabilities that were just the stuff of dreams in traditional environments. Even better, the business will find ways that it can turn IT into a business-facing powerhouse.

Foolish Assumptions

For this book, I'm assuming that you have at least a basic understanding of business and storage and the intersection of the two. If you're in a role that has anything to do with storage, virtualization, data center architecture, and business/technology decision-making then this book is for you.

Icons Used in This Book

Throughout this book, you will find a number of icons intended to help you better understand and remember key concepts.



I use this icon when you need to stop for a second and make sure you recall a key concept before forging ahead in a chapter.



Information pointed out with a Tip icon is something that may save you time, money, or both!



Make sure you heed the warnings presented in these sections so you don't fall victim to any storage traps.



If you're wondering how you connect the pieces of the storage puzzle together, you've come to the right place. In these technical sections, I peel back the covers a bit and show you how things work.

Beyond the Book

There's only so much I can cover on storage solutions in this short book! To learn even more about all-flash storage, check out these online guides and white papers:

- ✓ **Flash Storage Buyer's Guide:** <http://purefla.sh/2bjLDUJ>
- ✓ **MIT Technology Review: Flash Storage as a Strategic IT Asset:** <http://purefla.sh/2bsxXIIt>
- ✓ **ESG: Efficient Data Center Virtualization Requires All-flash Storage:** <http://purefla.sh/2btGw4t>
- ✓ **Forrester: The Total Economic Impact of Pure Storage FlashArray Storage Solutions for VDI:** <http://purefla.sh/2btGw4x>

You can also continue your flash storage education journey by visiting www.purestorage.com/flash and reading the resources available there.

Chapter 1

Containing Storage Costs

In This Chapter

- ▶ Adding up the TOC of traditional storage
- ▶ Bringing down your overall costs with all-flash arrays
- ▶ Using inclusive software licensing to bring down the TCO for your all-flash storage systems
- ▶ Knowing the importance of no-compromise all-flash storage

Creating business value is what drives the modern enterprise. Companies are always looking for ways they can save money in IT overhead so it can be invested toward revenue generating programs. IT and the data center — specifically storage — have traditionally resulted in a lot of money flowing out and into the hands of storage vendors and expensive storage staff.

Too few businesses understand the total cost of ownership (TCO) of their storage environments, making it difficult — if not impossible — to take proactive steps to control this cost. After all, if you don't realize you have a problem, you can't take steps to solve it. Further, after you've discovered that you do have a storage problem, what are some of the ways that the right all-flash storage solution can solve it?

Calculating the TCO of Traditional Storage

Identifying the problem is the first step in transforming your data center. In terms of storage, it's also the beginning of gaining an understanding about your full-loaded traditional

storage costs. The TCO for storage goes far beyond the purchase and maintenance invoices that you receive. Quite a few fixed and variable costs should be considered:

✔ Capital expenditures (CapEx)

- **Purchase price:** The price that you pay up front for your storage system is perhaps the most obvious part of the TCO. The hardware cost is almost always the biggest part of a storage acquisition.
- **Replacement cost:** Everything in your data center is on a replacement cycle. The time to decommission that expensive storage array will come. Your traditional storage vendor may recommend a full replacement of your equipment in order to stay current with technology and gain access to the latest and greatest technology.

✔ Operational expenditures (OpEx)

- **Maintenance:** From the beginning, your storage vendor expects you to pay an annual fee to maintain your equipment. It's basically a warranty plan that allows you to have the vendor fully support you in the event that something happens to your system. Choose the maintenance period to closely correspond to your life cycle replacement plans.
- **Power and cooling:** Many IT teams remain on either spinning hard disk drive arrays or on hybrid arrays that consist primarily of disk. With so many moving parts, disks generate a lot of heat and consume a ton of wattage. As a result, they require cooling systems to keep temperatures in check. The combination of increased cooling and the fact that spinning disks require more electricity mean that you're faced with increased power and cooling expenses.
- **Licensing:** With every vendor, you get an array. In order to go beyond that and add useful features, you may incur various types of licensing costs. For example, do you want data deduplication? What about cloning? Replication? Encryption? These features might be extra in the non-flash world. In fact, most of these typically require a premium cost. Licensing costs can add a significant amount of money to the overall storage cost of ownership.

- **Floor space:** With spinning disk systems, which remain in use in many organizations, attaining necessary performance levels can be pretty tough. To increase performance, storage administrators need to add more disks — spindles — to the mix. The decision to add more disks to address performance challenges results in data center racks that are full of whirring disks, spinning away, and taking up valuable — and expensive — data center rack space.
- **Complexity:** Many traditional storage systems are truly complex. The more complex they are, the harder they are to manage. As a result, your IT personnel costs related to storage management can be somewhat staggering. Even worse, this complexity can lead to slow downs for new business application deployments, increasing the costs for those projects as well. Complexity is an expensive characteristic for many storage environments, and while sometimes difficult to quantify, it has apparent and obvious cost implications.



All together, these items can create an iceberg-like storage ownership experience. The part that you see above the water that's obvious and visible may not seem intimidating. However, as you begin to peer below the surface, the hidden dangers begin to become more apparent. The true TCO for legacy storage becomes painfully evident.

Discovering the All-Flash Ingredients that Reduce TCO

You may not immediately associate all-flash storage with saving money. In the past, even mentioning the term “all-flash storage” may have resulted in tremors of fear because flash was once considered prohibitively expensive. Those days are long, long gone. Today's all-flash storage arrays are *far* more affordable than the original flash-based systems five or so years ago. Modern all-flash storage systems can be smart storage systems and can help organizations reduce their costs. This section explores how.

How data reduction amplifies available capacity

Even with all of the additional services that are often wrapped around storage these days, capacity remains the most important consideration for most organizations. And some interesting things are happening in the world of all-flash, thanks primarily to vast improvements in data reduction technologies.



Data reduction is an umbrella term that refers to features that can help organizations store more data into their existing storage capacity. You can effectively reduce the amount of data that you have to store through many different methods, each of which are discussed in depth in this book. For now, just consider things like data deduplication and compression as being the top two ways that data reduction and storage efficiency are achieved.



With data reduction, you get a multiplier effect that, in essence, amplifies the actual capacity of your storage. Data reduction technologies are measured based on a ratio that provides information about how much reduction you're getting. This ratio — represented as effective reduction: 1 — tells you how effective the array's reduction capabilities are based on your workloads.

For example, suppose you're seeing what would be termed as a "5:1" data reduction. With such a reduction factor, you're actually able to store the equivalent of five times as much information in every byte on your storage array, effectively increasing its capacity by 500 percent. Whereas data reduction is often an optional feature on legacy storage systems, modern smart storage systems integrate reduction technology right into the array, making it an inextricable part of the array's operation system.

Now, think about this from an economics perspective. For every dollar you're spending on *raw* storage capacity, you're actually getting \$5 worth of *effective* storage capacity. That's just one way that a smart storage system can begin to help you tame the wild beast that has become storage economics.

Facilities improvements: power and cooling

Facilities costs related to power and cooling are often a significant — and often overlooked — expense. Its benefits in an all-flash environment are derived thanks to physics as well as software features built into all-flash storage systems.

Physics? Yes, physics. Here's why: All of those moving parts in a traditional disk system use motors that consume a lot of power and generate a lot of heat, as well as a lot of noise. An array with a couple of dozen spinning disks can sound like a jet fighter taking flight! With an all-flash system, there are many fewer moving parts. While arrays themselves will still have fans to provide cooling, the disks themselves are motionless. The result is massively reduced electrical costs. Further, because *far* less heat is generated, far less cooling is needed, too. Your data center air conditioning systems will thank you after you deploy all-flash storage!



With flash storage, achieving rates such as 5:1 data reduction means that you need *much less* hardware than the past. Many legacy spinning disk-based data centers didn't have data reduction capabilities by default, so you only had raw disk space to work with. By enabling data reduction by default across an array, you can effectively reduce the amount of storage hardware you actually need, which further reduces power and cooling costs.

Data center storage footprint and cabling

Rack space can be really expensive. Square footage in any business is often a premium resource, especially in the data center. While virtualization helped to reduce the amount of physical equipment needed, it does have its limits. Over time, equipment space needs grow as new workloads are added.

Storage is one of the prime culprits, particularly as you consider a couple of key facts:

- ✓ Companies often add disk shelves when they need more performance.
- ✓ Many traditional storage arrays don't have data reduction capabilities.

After you move to an all-flash data center, both of these facts are no longer a concern. You won't likely need to add more shelves to add more performance. Flash storage systems effectively eliminate performance concerns. On the capacity front, flash disks are increasing in size all the time. Between that and the fact that and data reduction features, you don't need as much physical hardware.



As you reduce the number of racks and appliances you have in the data center, you can also reduce the number of cables that you need to run to tether everything together, which simplifies the data center. Fewer cables translates to less air movement restriction, which can also reduce cooling costs.

Eliminating Storage Software Licensing Headaches and Costs

The legacy storage purchasing process is arduous and difficult. At first glance, you might simply look at array pricing, but you soon discover many optional add-on features and software that aren't included in the base configuration, not to mention lengthy maintenance and support contracts.

And what if you add an operating system on your array? That'll be extra. You want data reduction in the form of deduplication and/or compression? Extra. You want replication? Extra. If you decide to add on optional manufacturer software features, then comes the fun part: actually getting the license keys to unlock those features. Even after you pay the legacy vendor, getting license keys can sometimes be a challenge in itself.

Now imagine a world in which you pay once and get everything you need. Sounds like a dream, doesn't it? It's not! Today, a few of the leading all-flash vendors include software for free with the array purchase. Some of these features, such as deduplication and compression, are critically important to the flash economic picture, and attempting to decouple the component would result in a less-than-stellar experience.



Bundling all the necessary and optional software components right into the base system means that you have less hassle and less expense. That's a winning scenario.

The Importance of Smart Storage

For too long, you worried about your storage in some way. You were uncertain about its capability to keep up with your workload demands and how to tier them appropriately. You were concerned whether the arrays would suddenly become unsupported. You weren't sure if you had the budget to fund software feature upgrades to keep up with evolving business requirements. It's time to change the paradigm.



With smart storage, those worries go away, particularly when the storage purchase is coupled with a free software license and includes ongoing support options. In some cases, this also means that you have the peace of mind of not needing to rebuy your storage for up to a decade.

10 All-Flash Data Centers For Dummies, Pure Storage Edition _____

Chapter 2

Reaching Simplicity and Manageability

.....

In This Chapter

- ▶ Learning why simplicity is important in the data center and in storage
 - ▶ Finding out the ways by which a simpler data center translates to a lower TCO
 - ▶ Debunking old flash myths
-

IT teams are under constant pressure to reinvent themselves as new technologies hit the market and as business requirements evolve. Today, the world of technology is undergoing cataclysmic change as emerging options provide more choice and flexibility for the business. No longer can IT professionals afford to immerse themselves in complicated cabling nests or deploy infrastructure that requires months of training to manage.

In this chapter, you discover why simplicity has become a meta-trend across all aspects of the data center, including storage and how simplicity in the storage silo can have positive impacts on many aspects of the business.

Why Simplicity is The New Black: Scaling the Business

There was a day when deploying data center products that had complicated administrative interfaces might have been considered a positive, at least by IT. After all, if it's complicated, it must be powerful . . . or something like that.

No longer.

Today, smart IT buyers look for core storage features such as

- ✔ Simplicity
- ✔ Ease of use
- ✔ Streamlined administrative experience
- ✔ Automation
- ✔ Simple procurement process

Organizations of all sizes are proactively seeking ways that can make IT simpler to manage. Business and IT leaders want to ensure that budgets and resources are invested in business critical areas. And all new technology enhancements are easy to deploy and leverage the latest technology to streamline efficiencies in the data center.

Traditionally, storage has been among the most expensive and complex resources managed in the data center. It's no surprise that all-flash, smart storage solutions have become mainstream and effectively eliminate storage performance challenges by offering breakthroughs in performance, scalability, simplicity and agility — not to mention that they're simply easier to procure, manage, and maintain.



Only through simplicity can an organization easily scale its business. After all, if you're constantly reinventing the wheel inside the business, your focus isn't on the long-term.

How Does Storage Simplicity Manifest in an Organization?

Traditionally, storage is one of the most complex parts in the data center; therefore, bringing simplicity to storage management can offer a number of important outcomes for your company.

Reducing procurement complexity

In terms of cost, storage simplicity has a number of outcomes. First, simplicity reduces the amount of time you spend procuring storage. Here's proof: Were you to compare invoices for storage purchases from a purpose-built all-flash vendor and a traditional storage vendor (even legacy vendors with all-flash options), the non-legacy vendor invoice will probably show just two to three items for a purchase. On the other hand, legacy vendors may actually have pages upon pages of options, products, and services. So, right from inception, a simple all-flash storage system begins to save you money in terms of procurement complexity and time.

Repurposing IT resources

Everyone wants to do more with less . . . or, at least, with the same. Whenever you can reduce the overhead associated with managing a resource, such as storage, the more you can pay attention to other areas of the business. Any reduction of storage-induced overhead helps your team better manage valuable IT resources. Of course, actual time and resource savings vary based on your organization's unique requirements.

Reducing the data center maintenance burden

Beyond simply being able to redirect IT staff efforts, the simpler the data center, the more quickly you can deploy new workloads. The reason is simple: If you have less complexity in the data center, it's less likely that something will go wrong as you try to deploy new workloads. In other words, simplicity in the data center reduces the overall data center maintenance burden and can directly benefit the business by making it easier and faster to deploy new business workloads.



Further, a maintenance-related situation that you should seek to avoid is one that's really common: unexpected maintenance costs. When you first buy that shiny new legacy array, you may be delighted to find that the maintenance costs are downright reasonable.

Over time, though, insidiousness begins to sneak into the storage maintenance equation. Vendors increase the cost of maintenance as the likelihood for catastrophic breakdown increases. Eventually, you have to make a decision around completely replacing your storage because maintenance costs get so high. In other words, you have to rebuy your storage every few years and then pay an increasing amount on an annual basis to support and maintain it.

In order to combat this phenomenon, Pure Storage has developed a new program, known as Evergreen Storage. This innovative program is designed to extend the life of the storage deployment to more than *ten years*. Pure Storage seeks to accomplish this by enabling IT organizations to augment existing FlashArray//m deployments with upgraded processing, memory, and solid-state storage as new innovations emerge, all while the data is online and available. The net result is intended to allow existing configurations to scale performance and capacity over time as more advanced and less expensive technologies emerge.

For many organizations, the combination of Pure Storage's architecture and Evergreen Storage provides the benefits that scale-out provided to spinning hard-drive environments in an efficient manner, without locking capacity and performance scaling together.

Tuning Flash Storage for Application Workloads

If you're buying an all-flash array, it's likely that you want to maximize the investment by ensuring that your mission-critical business workloads operate at peak efficiency. You don't want to make that investment only to find that the storage isn't performing as you expect.

Even with some flash solutions, you may still need to tune your application to achieve optimal results, but that's often due to inflexibility in the storage design. For example, for flash platforms with fixed block sizes, applications may need to be exported and imported to align to the I/O size; this process is a common recommendation with OLTP databases, VDI, and virtual infrastructures.

To that end, with some all-flash systems, customers need to spend significant time tuning the array to match performance and operational characteristics of their workloads. This is generally due to some deficiency in the platform. Or, it could be that the flash vendor has taken a legacy disk-based array and attempted to retrofit it to support all-flash. This method rarely yields good results; an array that's built with native support for all-flash will almost always outperform retrofitted systems.



Here are some other considerations. If you need to load balance your application across volumes and/or controller nodes, you may find that some flash solutions limit performance on a per-volume or per-node basis. You may also find that the array requires considerations for application data layout both at application inception and in the future as an application grows.

Busting Flash Reliability Myths

Flash fails. That is simply a fact. But so do hard drives. So does your car's engine. So does everything else, at least eventually.

In the very early days of the modern flash revolution, flash was branded as unreliable due to the fact that it seemed to fail more quickly than spinning disk. While the problems inherent in early flash have been solved over time, these early failures remain in people's minds and, to them, serve as cautionary tales about moving to an all-flash data center.

If you've been around flash storage even a little bit, you've heard that there are different kinds of flash storage media. Today, the most common kind of flash media available on the market is known as *multi-level cell* (MLC). MLC brought to the flash market a steep drop in price as it was capable of storing multiple bits per flash cell. Prior to the introduction of MLC, single-level cell (SLC) ruled the roost. However, it was prohibitively expensive because it could store just one bit of data per flash cell.

Figure 2-1 provides you with a look at the different kinds of flash that are on the market today.

	SLC	eMLC	cMLC	TLC
Flash Cell P/E Cycles	100,000	10,000	3,000	300
Performance	Highest	High	High	Medium
Price per GB	Highest	High	Moderate	Low
Bit Error Rate	Low	Low	Low	High

Figure 2-1: The types of flash available on the market today.

You may be wondering why I'm including this technical material here. There's a darn good reason. In Figure 2-1, you will note that there are two different kinds of MLC media. One is called enterprise MLC (eMLC), and the other is known as consumer MLC (cMLC). It would seem intuitive that you'd insist on eMLC for your enterprise workloads, right?

Well . . . no.

You see, the terms eMLC and cMLC were dreamt up by marketing departments seeking to differentiate the two media types. While it's true that budget-friendly cMLC cells may only survive 3,000 P/E cycles (a P/E cycle is initiated whenever a cell on flash media needs to be overwritten), that doesn't necessarily mean that you need to jump to expensive eMLC media, which can survive 10,000 P/E cycles.

It all comes down to how well the flash-array vendor has implemented what are known as *write avoidance* techniques. Write avoidance is exactly what it sounds like. If you can avoid having to write data, you reduce the wear and tear on the flash media.

Now, this doesn't mean that the array is going to stop saving your data just to extend the life of the individual disks. However, what it will do is leverage software features, such as data deduplication and compression, to reduce the number of times that data has to be written to the array. Through a combination of techniques, an all-flash vendor can massively reduce the number of times that data has to be written to the array, which significantly extends the life of that flash media.

Chapter 3

Supporting Virtualization and Consolidation

In This Chapter

- ▶ Learning why disk-based virtual environments can cost you money
- ▶ Solving vexing virtualization and consolidation performance challenges

The modern data center has its roots firmly planted in the outcomes of trends that began in early 2000s. Back then, x86-based server virtualization was brand new — a trend that gained traction as organizations realized the kinds of benefits that they could glean from virtualization technology.



Virtualization benefits include the following:

- ✔ **Reduced costs:** Virtualization consolidated many workloads that were formerly housed on individual servers and helped companies actively reduce the amount of hardware deployed in the data center.
- ✔ **Easier management:** With fewer physical servers comes less overall management. Virtualization initially collapsed workloads onto less hardware, and, more recently, networking capabilities have also become increasingly folded into the virtualization environment.
- ✔ **Workload mobility:** The ability to shift workloads from on-premises to the public cloud and between data centers in the private cloud is increasingly important and is a well-supported virtualization scenario.

Identifying Key Storage Performance Challenges

Unfortunately, not everything is perfect when it comes to virtualization. In fact, when it comes to storage, there are some key challenges that must be overcome in order for the full potential benefits from virtualization can be realized.

Performance concerns

The old days of virtualization storage were based on spinning disk, especially with I/O-intensive workloads such as VDI or analytics. In these cases, obtaining sufficient performance from spinning disk systems is, suffice it to say, a challenge. To get more performance in their disk-based storage environments, companies have leveraged various techniques:

- ✔ **Short stroking:** If you analyze the physics of a spinning disk, you will find that the outer portions of the platter spin faster than the inner portions. As a result, data stored at the outer edges of a disk can be retrieved and written more quickly than data in the inner portion. To leverage this physics-based reality, some have taken to configuring their storage systems to use *only* those fast portions of the disk. On the plus side, this results in improved performance, with some metrics claiming a more than 50 percent performance increase. However, this technique comes at the expense of capacity. By using only a part of the disk, you're effectively forced to deploy additional disk shelves to meet your organization's capacity needs.
- ✔ **Throwing hardware at the problem:** More "spindles" means more Input/Output Operations Per Second (IOPS). Every disk you add to a storage array adds more IOPS. As performance became a challenge, people have simply added more disks, even if they didn't need the capacity.



These approaches are really expensive and complicated. When you throw hardware at the problem, you're increasing all your costs — rack space, power, and cooling — not to mention the fact that you might be buying storage capacity that you don't need. Of course, in the pre-flash days, these were the only options available, so it makes sense that they were pursued.

The rise of server-side storage solutions — both hardware and software — is intended to help solve storage performance issues. These products create server-side caches that reside on either PCIe-based flash storage or on virtual disks created from RAM that has been carved away for this purpose.



These approaches consume valuable server resources, including RAM and CPU. The resources that you dedicate to solving storage performance problems can't be used for workloads — you may have to deploy additional hosts in order to run all the workloads that you need.

Lack of budget

You'll need more capacity at some point. Being able to *buy* additional capacity might be really expensive, for a variety of reasons. For example, because disk-based systems (and even some all-flash arrays) don't always include data reduction technologies, such as deduplication and compression, you're paying full-fare for all your storage capacity. Even those that provide data reduction technologies do so with varying degrees of efficiency. These factors can mean that upgrading your disk-based SAN to increase capacity and performance can be prohibitively expensive.

Purchasing additional infrastructure

Budgets are always challenging, but it's often *the* driver for many organizations. Imagine needing to add a bunch of disks to your environment just to improve performance. By doing so, you're also increasing your power and cooling budget challenges. After all, those disks require power to spin and cooling to reduce the heat they generate.

By forcing the need to increase storage acquisition costs as well as massively increasing power and cooling costs, you might need to increase your cooling capacity, which requires additional infrastructure. Some storage arrays scale performance and capacity independently; others require you to scale in fixed, hardware — defined units (think appliances or hyper converged infrastructure) that require additional power supplies and multiple network I/O ports. Lastly, server based

scaling requires redundancy overhead to protect data and provide high availability (HA).

Complexity

Far too many legacy storage systems require deep skills and expertise to implement, manage, and maintain. With virtualization comes an expected level of simplicity, but this simplicity outcome is only recently beginning to encompass the storage resource.

Many legacy disk-based storage systems suffer from complexity issues that force administrators to read and memorize arcane tomes of work in order to perform even basic operations. As a result, organizations need to maintain multiple teams of expensive IT professionals. One team manages the hypervisor and workload side of the equation and another team manages the storage. This isn't desirable as it results — at least some of the time — in turf battles, which can add even more delay to getting new services deployed.

Solving Virtualization and Consolidation Challenges

Flash storage can help you solve even your most pressing virtualization and consolidation challenges. Even better, flash storage can solve your performance challenges while saving you money (and help you make the business more agile).

Reduced latency

Of every metric that you can possibly use to measure storage performance, none is more important than latency. *Latency* is the amount of time that it takes for an application to receive a response from storage after a request has been made.

Low latency equates to good performance. Every other metric — IOPS and throughput — are wrapped up in the latency metric. If you're not getting good IOPS, for example, latency increases because it takes more time for the media to complete operations.



With flash, though, IOPS isn't a problem. Even as you consider different storage block sizes, all flash arrays always beat spinning disks in terms of raw IOPS. The same goes for throughput, although things are a bit different. With flash, raw media (disks) are almost always going to provide a lot of IOPS. However, as you transfer that data from the disk to the application, it goes through controllers and other constructs. If you're working atop a storage architecture that's been retrofitted to support flash — originally designed for disk — your throughput may suffer, and you won't get all of the speed you could if you were using a native-flash array.

If you're using a native-flash array, you get a very low latency (sub 1ms is best in class), which translates to much better business application performance. It also improves the overall performance of your virtual environment.

Liberating server resources to improve VM density

Some products out there can be installed on virtualization hosts to help boost storage performance. These products, while innovative, can have a severe impact on how dense you can make your virtual environment. Although you may solve the storage performance problem, you're doing so by forcing yourself to deploy even more hosts to support the same number of workloads. As you consider server, hypervisor, and additional power and cooling costs for these servers, the cost can become prohibitive.

With a smart storage system built to natively support all-flash, you can eliminate the need for these “server-side tricks” and allow your storage system to function how it's supposed to function. You can also increase the number of virtual machines that you run on those hosts.

Deduplication

Sometimes, consolidation challenges aren't about performance. Instead, they're about the sheer volume of data to be considered and supported. However, as you pull back the curtain, an interesting fact emerges: There's a whole lot of duplication in your storage system.

Suppose you have 250 virtual machines. Is each one of those VMs running a fully unique operating system? Probably not. Instead, you might have 250 Windows Server 2012 R2 virtual servers deployed. While each one of those servers might serve a different purpose, the fact is that you've installed the same operating system 250 times, so a great deal of commonality exists across those virtual machines.

Deduplication has some serious capacity-saving potential. Better yet, it can even result in massively improved performance. Here's how: Writing data to flash storage is pretty fast, but that speed is impacted as you put cells through the program/erase cycle, which adds latency to a write operation. What if you could *not* write the data at all? That would be pretty darn fast!

With deduplication, that's exactly what takes place. If the array determines that a block of data already exists on the system, a new request to write the same data will simply be rejected. You find out more about deduplication in Chapter 4.

Management integration with the hypervisor

For a long time, the hypervisor and the storage environment were separately managed entities, each featuring its own administrative console and constructs. Administrators, therefore, had to learn two tools. However, while having to learn two tools isn't necessarily terrible, what *is* terrible is having to go to two places every time you work on a virtual machine.

Fortunately, times are changing. As storage vendors more fully integrate their solution management with the hypervisor management tools, storage administration is becoming far easier than it used to be. Smart storage systems often provide management tools that can help ease storage management activities and free up resources for other projects within the data center. Additionally, integration into hypervisor management tools makes it easier to adapt to changing performance demands of new workloads.

Chapter 4

Predicting and Modeling Future Storage Consumption and Costs

In This Chapter

- ▶ Finding out how long-term storage economics are changing
 - ▶ Discovering how to eliminate storage-induced budget spikes
 - ▶ Meeting your storage goals while reducing costs
-

If you deploy the right all-flash storage solution, you can tell your company exactly how much its storage will cost and you can save money in the process. In this chapter, you discover the key economic factors that you need to know to be successful in your storage endeavors.

Understand Long-Term Storage Economics

I discuss in Chapter 1 the various components that comprise your storage environment's total cost of ownership (TCO). Flip back there for a quick recap if you need one. These components form the basis for your long-term storage economics. As you review this information over a 10-year period, the numbers might be downright frightening. During that time, you may buy storage and then replace it at least once. You need to keep it maintained. If you buy the wrong system, you need to budget for licensing and for complexity.

All this talk of replacement and cost can be a bit scary, but don't be afraid. There's a solution. Smart storage systems today can drive major costs out of the storage equation and make your matrix look more like what's shown in Figure 4-1.

CapEx	OpEx
Purchase price	<p>Maintenance: Predictable and fair</p> <p>Power and cooling: Massively reduced</p> <p>Floor space: More storage density equates to less floor space, which equates to lower costs</p> <p>Periodic technology refresh: Never again re-purchase or pay for upgrade costs for periodic hardware refreshes; in fact, make performance, scale, and feature upgrades automatic</p> <p>No capacity surprises: Guaranteed effective capacity, based on your initial storage workloads; protects your storage investment</p>

Figure 4-1: A storage environment build on top of evergreen storage principles.

Typically, a *budget spike* occurs whenever a purchase needs to be made that's out of the ordinary. Many capital expenditures result in budget spikes. In the world of IT, these budget spikes happen every time you replace a resource as it hits the end of its life cycle — around every three to five years. This is one of the main reasons that storage is considered to be the most expensive resource in the data center.

Even worse, when you replace your storage, you're also subjecting the company to some level of disruption. You need to schedule downtime for applications so that you can move them to the new storage.

CFOs *hate* budget spikes like the one shown in Figure 4-2.




Figure 4-2: An exaggerated budget spike . . . for effect.

What if you could eradicate the budget spike with regard to storage once and for all? What if, after you've purchased your storage, you could keep it for ten years or more and be able to predict your costs?



You can. One storage vendor, Pure Storage, has made available what they call Evergreen Storage. With the Evergreen Storage guarantee, a number of very cool things start to happen, many of which I describe in this chapter. With Evergreen Storage, you never have to rebuy a TB you already own. Evergreen Storage behaves like SaaS and the cloud. You can deploy it once and keep expanding and improving it for ten years or more, all without any downtime, performance impact, or data migrations. With Pure's Evergreen model of technology and biz practices enabling a system to last ten years or more, customers can avoid two to three upgrades during a ten-year life cycle.

Discover Inline Deduplication's Cost Benefits

There are two primary ways by which you can achieve data deduplication goals. The first method, known as *post process deduplication*, is kind of like a backup window. In essence, all data is written to the disk, and then, on a schedule, the array deduplicates it. However, scheduling data deduplication may

cause a loss in capacity savings if it's not timed correctly. For example, some data may be locked in snapshots and replication processes. Further, you may end up in a situation in which your deduplication window is too short, and you either can't reduce all of the data or the process runs into working hours and potentially impacts production.



With an all-flash system, *inline deduplication* is not only possible due to the performance of flash and new architectures designed for it, but also inline deduplication is critical in raising the reliability of flash by avoiding writes and eliminating PE cycles for data that already exists. Inline deduplication eliminates the challenges historically associated with disk storage.

Guaranteed data reduction

Not all data reduction is created equal. So, it should come as no surprise that, as you review your all-flash storage options, you find that different arrays provide different outcomes with regard to data reduction. There are a number of different reasons for this.

Deduplication block size has a lot to do with things. The block size granularity at which they analyze data with fall into one of two categories based on the storage platform's back end architecture — they're either fixed or variable in length. The smaller the granularity of the block size used to analyze data, the more redundancy and capacity savings will be returned.



Variable length deduplication provides the greatest level of granularity, allowing redundancy that would normally not be identified with a fixed block implementation to be identified. This model requires more CPU cycles but returns the greatest possible data reduction.

Fixed length deduplication limits the granularity of the block size to match that of the back end. This form of implementation optimizes performance by limiting the number of CPU cycles dedicated to identifying redundancy at the expense of lower reduction levels.



For the best results, find a storage solution that provides variable length deduplication. Of course, there's more to data reduction than just deduplication. Data compression is an important component in the data reduction story. But, there are different kinds of compression algorithms out there, too. Compression provides a varying range of data reduction based on the data set. Some data responds better to one form than another and some respond well to a combination of algorithms. Flash storage that offers multiple data compression algorithms can be relied upon to provide a quantifiable advantage in data reduction over those with only a single form.

Most all-flash vendors will claim some level of effective capacity. Effective capacity is defined as the capacity you're able to achieve with an array once you take into consideration all of the data reduction features. But, not all workloads are created equal and neither are all arrays.



To that end, as a part of its Evergreen Storage guarantee, Pure Storage, based on the kinds of workloads that you're operating, guarantees a certain level of effective capacity for up to six months after you buy your array. With such a guarantee, an all-flash storage solution is more than economically viable.

Predictable performance

Capacity and economics are important, but the entire point of storage is to support workloads. Perhaps the biggest benefit that an all-flash storage system brings to the table is the ability to run workloads with ongoing predictable performance. From the perspective of your users, predictable performance is, perhaps, the best possible outcome . . . as long as it's also fast. How fast? Sub 1ms latency enables consistent, fast performance for many workloads, including database, virtualization, VDI, ERP, HER, and so on. Compare that to disk technology, for which single-digit ms response times are difficult to achieve even in the best of circumstances.

It also means that, as you add more capacity to your environment — for example, perhaps you need more capacity in your second year of deployment — you can plan on continuing to increase overall performance. As you expand with additional nodes, performance will scale linearly.



As you consider all-flash storage options, you also need to be able to survive failures. Any failures that you endure shouldn't be debilitating. With the right smart storage solution, if you've suffered, for example, a controller loss, you should be able to maintain full performance. The same goes if you have a controller offline for an upgrade.

With some arrays, suppose there is a 50 percent performance degradation during upgrades under load. In such instances, customers are likely to leave that array partially unloaded in order to lessen the impact of upgrades. With a properly designed storage environment, there is never a performance degradation during upgrades, or even failures — and you get full value for your investment.

Chapter 5

Achieving Scalability and Agility

In This Chapter

- ▶ Bringing simplicity and agility to the storage resource
 - ▶ Taming the resources in your data center, including storage
 - ▶ Growing your business with APIs and predictive performance
-

Your budgets are often staying level or, worse, shrinking, even while you're being asked to continually ramp up the level of service that IT is offering the business. And at the same time that the business is demanding more services and tighter service level agreements (SLAs), it's also demanding that IT protect the business to a greater extent than ever before. This adds yet more burden to IT.

So, how do you accomplish the seemingly incompatible goals of helping to grow the business while also keeping the budget in line and protecting the business from disaster? This chapter helps you figure that out.

Using Simplicity and Operational Effectiveness for Scalability and Agility

You must find ways to make your budget dollars stretch further and your technology purchases have more impact. This is where increasing operational effectiveness and introducing simplicity into the data center come into play. You can take

a number of storage-related steps to start taking control of costs and operations.



Help your business forge ahead by providing a foundation that enables easy management and scale while allowing the business to remain flexible and agile.

Consolidate storage-related services

Today's storage systems do way more than just provide a place for you to store your data. In fact, with the right *data services*, your storage can be a critical component of your journey to simplicity. This is a major paradigm shift! Storage used to be the resource that held people back; today, the right solution can help propel the business forward.

Are you running separate appliances for storage, data reduction, and disaster recovery? Stop! Consider instead the adoption of a storage solution that handles all of the functions natively. By doing so, you immediately remove some of the complexity from the data center equation. You also remove multiple maintenance figures from the budget equation, and you reduce the number of points of support contacts that you have to manage.

Fewer touch points and fewer budget lines make your IT operation a lot simpler. This simplicity can free up budget that you can use for other business-facing needs. Further, as you need to grow the environment, it's easier to just add more of these modern, efficient systems.

Adopt converged infrastructure

There's more than just storage in your data center. All these additional resources mean that the complexity levels in your data center just continue to rise. Making everything work together can often be time-consuming. Moreover, integrating solutions from different sources can introduce a lot of delay when it comes to deploying new business applications. Agile it is not. Therefore, it stands to reason that the business is less agile due to this complexity.

Of course, with all of this time and complexity comes direct and indirect expenses. *Direct expenses* include the amount of staff time that you expend on these efforts. *Indirect expenses* include the negative return on investment (ROI) during the period that the equipment is in staging and the opportunity cost for all of the above.



Converged infrastructure — a loose conglomeration of servers, storage, fabric, and a hypervisor — was developed to solve some of these challenges. Unfortunately, many existing converged infrastructure solutions are based on old, slow, inefficient disk storage or on unreliable and expensive flash systems. The resulting architectures are too expensive, perform poorly, or have significant downtime windows — making them unsuitable for the consolidation of business-critical applications and workloads. That doesn't mean that convergence is the wrong answer. It just means that you need a converged infrastructure solution that has the right characteristics, such as

- ✓ High levels of performance and scalability
- ✓ Simple to operate and doesn't add unnecessary complexity to the data center environment
- ✓ Enterprise-class equipment to enable dependable support for all applications
- ✓ Economically viable and able to help you reduce your overall infrastructure, power and cooling costs while also providing all of the storage benefits



By deploying the right converged infrastructure solution, you make it simple to scale the data center environment without having to worry about extended deployment times and negative return on investment (ROI). Pure Storage FlashStack with Cisco is a tested, pre-validated converged infrastructure solution that dramatically shortens time to value for mission critical workloads of all sizes.

Manage with the right tools

As you're evaluating storage solutions and as it pertains to scalability and agility, you discover that you need better management tools to achieve your goals. You need a storage administrative interface that works from anywhere, providing you with real-time, at-a-glance information about the details

that matter, such as current storage capacity, level of experience data reduction, read and write latency, and throughput. You shouldn't need to go digging for these important details. Figure 5-1 provides an example of an administrative interface that gets the job done from any device. Information at your fingertips aids data-driven decision making around the need to scale very easy.




Figure 5-1: Pure Storage Pure1 administrative interface provides real-time at-a-glance details.

Stop worrying about upgrades

Many companies discover that they've hit an expansion wall with some of their storage solutions. Or, they discover that they're now running storage systems that are end of life or otherwise unsupported. These kinds of scenarios can introduce a ton of unanticipated cost into the storage equation and significantly extend deployment times.



Upgrading shouldn't be a worrisome affair. Don't lock yourself into a situation where you're forced to rip and replace something. That absolutely *kills* your business's agility. Pick a storage solution that takes the worry out of upgrading whether it's during a technology upgrade or for maintenance activities. Get the ability to upgrade capacity and performance with no

downtime required and maintain a storage environment that can last upward of ten years. Then, storage is never a stumbling block.

Driving Scale with Predictive Application Performance

Predictability is often more important than raw performance. After all, predictability is what drives certain guarantees made to the business, such as SLAs. Performance predictability comes in two flavors. First, it means that you can guarantee certain levels of IOPS, throughput, and latency. However, “performance” often refers to availability, too. In the case of storage, both definitions of the term are really important to keep in mind.

On the performance front, your storage system shouldn’t provide an inconsistent experience. The kind of performance you see in day-to-day operations shouldn’t be impacted if you, for example, lose a controller or are performing an upgrade. Here’s why: If the applications are mission-critical and require guaranteed performance and if the loss of a single controller would cut performance in half, you might only load up the array with one-half of the overall application load and buy a second array for the other half.

That isn’t a recipe for good scalability! Only when you’re able to guarantee that performance levels will always remain consistent and predictable can you comfortably consolidate applications to less hardware to achieve both the ability to easily and predictably scale and to reduce your costs.



Having good monitoring ability helps with predictability, but so does a system that can deliver high availability. Look for a system with no single point of failure to ensure resilience.

Leveraging APIs for Automation

Automation is a trend that’s popular because they massively reduce what businesses spend on human labor. In the spirit of doing more with less, automation is also a key need in the data center. Modern data center equipment manufacturers are including really powerful application programming

interfaces (APIs), which enable IT staff members to script against the hardware. In some circles, the inclusion of APIs is said to create an “infrastructure as code” environment in which everything is easily programmable.

By having the ability to script operations, organization can automate routine tasks, which significantly decreases the number of times that a human has to physically touch the administrative interface. A secondary impact is that there’s often a decrease in human error — the less that humans touch something, the less likely they’ll break something.

Supporting Data Growth

The primary “scale” consideration when it comes to storage is often capacity. With flash, many are concerned that flash media can’t hold as much as spinning disk. Not so fast. In reality, there are some really interesting things happening that directly contradict this myth.

First, you have to consider effective capacity, not raw capacity. *Effective capacity* is what you get after you apply your data reduction ratio. For example, suppose you have a 20TB array and you’re getting 5:1 reduction. In reality, you have an array with 100TB of effective capacity. Of course, different kinds of workloads will yield different kinds of reduction results.

Today, most all-flash storage vendors sell arrays that periodically dial back to the vendor mothership and upload all kinds of telemetry. This information doesn’t include your sensitive data, but it does include information about the overall health of your array and hints at what applications you’re running. Armed with that information and after analyzing all of this data, these vendors can begin to provide you with some guidance as to what you should experience with your own application mix.

With the right smart storage solution, as your business grows and needs more capacity, and as technology continues to advance, don’t be left behind. For example, with its Evergreen Storage program, Pure Storage customers can continually leverage new advances in storage technology and capacity without having to re-buy what they’ve already paid for once. When it comes to supporting data growth as the business scales and as the business needs to remain agile, it doesn’t get much better than that.

Chapter 6

Addressing the Cloud

In This Chapter

- ▶ Discovering two disruptors in the data center
 - ▶ Understanding the differences in cloud
 - ▶ Finding the key ingredients for enterprise cloud
-

The cloud. It's a thing. And it's here to stay. There are a lot of ways to think about the cloud and how it impacts your data center and your IT organization. In this chapter, you explore the cloud in depth and look at the reasons that an all-flash data center makes a great foundation for your modern data center.

Looking at Flash and Cloud as Key Data Center Disruptors

Walk around any technology conference or show today and the words “flash” and “cloud” are thrown around like candy — and for good reason. These technologies have proven to be among the most disruptive since the advent of x86-based virtualization, and they show no signs of slowing down.

Flash has turned the storage world upside down. No longer are companies relegated to “throwing hardware” at storage challenges. Today, companies can get all the performance they need by deploying an all-flash data center with smart storage. Better yet, this performance doesn't need to come with a capacity compromise (check out Chapter 5 for more info).

A technology that has also disrupted the IT market is the cloud. You've heard about all kinds of different cloud environments, including public cloud, hybrid cloud, private cloud, software-as-a-service (SaaS), infrastructure-as-a-service (IaaS) and more. No longer is the business forced to deploy new applications in a local data center; instead, business units can go out to a provider on their own and stand up new services without even involving IT.



This is both a risk and an opportunity for the business. By going around local IT, business units may put the environment at security risk, and it becomes far more difficult to achieve economies of scale by centralizing purchasing power. However, in some cases, enterprise IT has earned a reputation for being slow — a trait that isn't compatible with propelling the business ahead.

Some folks in IT have looked at the various cloud providers with trepidation because it poses a risk to business as usual in IT. Rather than looking at the cloud as competition, IT needs to take a consultative role with the business to help point them in the right direction with regard to cloud and to corral the various services in order to protect the business as much as possible.



For those that are truly terrified that the cloud represents an existential risk, it's time to look inward. What is your local IT department not providing that's pushing the business to look at cloud? After you discover those traits, you need to bring them in-house and change your company's computing paradigm.

Understanding Cloud Characteristics versus Destination

Consider the word *cloud* for a moment. For many, when this word is used, they immediately think of public cloud services and the companies that run them. For others, their thoughts turn immediately to the private cloud they're attempting to build in the data center.

The private cloud fallacy

Most of you reading this book have heavily virtualized your data centers. There are those that believe that achieving high levels of virtualization means that they've successfully deployed their very own private cloud.

Nothing could be further from the truth. Cloud is about characteristics, not a destination. One of the characteristics of a cloud environment is virtualization or containerization, which, under the hood, is still a form of virtualization. However, in order to be

considered a cloud-like environment, the data center must have some additional characteristics, including some level of automation and user self-service, some level of resource pooling, enabling charge-back (this doesn't mean that it's used, but it should be possible), and enabling rapid elasticity, just like a public cloud.

Without most or all of these characteristics, a heavily virtualized data center is just that — heavily virtualized. It's not a private cloud.

In most cases, people's thoughts immediately jump to the workload destination when, in reality, they should be thinking about the characteristics that define the term instead. By thinking of cloud as a set of characteristics rather than a destination, some interesting things happen. First off, the "public cloud versus private cloud" debate goes out the window, as it should. Frankly, neither is going to *win* in the market. The future is firmly planted around *hybrid cloud*, which is a combination of public and private. There's a good reason for this: People aren't necessarily looking to just chuck their workloads to a cloud for the heck of it. They're looking for the kinds of outcomes that are possible only with a cloud-like environment.

Looking to All-Flash Data Centers as the Foundation for Enterprise Cloud Strategies

With an understanding that *cloud* is a set of characteristics rather than a destination, in this section, you take a closer look at these characteristics as they pertain to the storage layer.

Simple scaling

A key ingredient for cloud success is the ability to quickly and easily scale the environment from two different perspectives.



As overall business demand dictates, you should be able to easily add capacity to the data center. This capacity can come from physically adding hardware to enabling data reduction services such as deduplication and compression to increase the overall effective capacity of the environment. An all-flash solution purpose-built to support flash will provide this kind of scalability.

Predictable performance

Deploying storage that provides you with predictable levels of performance at all times is highly important. This means that the storage layer needs to maintain consistent levels of performance even when a controller fails or when you've taken a controller offline to perform an upgrade.

Just as important, as you scale the environment, you should not suffer performance loss. The performance increase should be linear; so, if you add an additional array to your existing array, you should double performance.

With the predictable performance from a modern all-flash storage system, you can begin to make *real* service level agreements to the business, and it can fully rely on the services you're providing.

Always-on

Downtime isn't acceptable. Beyond providing predictable levels of performance, routine operations, such as firmware upgrades, capacity expansion, and even controller replacements, shouldn't result in downtime. With many legacy disk-based arrays, different volumes and sets of disks are pinned to a specific controller. Although a secondary controller can take over the burden if the primary controller fails, the end result is impaired performance, which is about as severe as an outage for some companies.



Some of today's modern flash-first architectures avoid these kinds of scenarios, but make sure you delve into the details in order to choose the right one. They're always on, even when a component is down for servicing and without impact to user or application performance. That is the kind of architecture that you need for your 24/7 cloud.

Easy management from anywhere

The chances are really good that you'll need to manage your storage environment at some point even if you can't get to the office or to a PC. Make sure you can use your laptop, tablet, or your smartphone to manage infrastructure remotely when it's not convenient to be in the office.



If you're using a storage management tool that isn't driven by HTML5 and requires <shudder> Java or some other monolithic legacy construct, you're being actively held back by your storage vendor. You're giving up invaluable family time to find a PC when something goes wrong. And, worst of all, reacting to a storage issue can take a lot of time as you attempt to find a place from which to administer the environment.



A key ingredient for success in your enterprise cloud is the ability to easily manage your storage from any location and from any device. Look for this capability in smart storage systems.

OpEx cost models

In a private data center, there will almost always be a combination of CapEx and OpEx expenses that need to be considered. The data center economic model is changing dramatically, primarily due to cloud and shifting budget priorities.

With an all-flash data center based on modern technology, you can begin to provide cloud-like services to individual business units using a fractional consumption OpEx-focused economic model. Your business can gain the internal economic benefits of cloud computing.

Integration capabilities

You already know that storage is but one of many resources in the data center. You also have compute, the hypervisor, and the fabric to think about. You need a storage system that can integrate neatly and completely into a converged infrastructure stack. Moreover, you need a storage system that includes an *intentionally designed* converged infrastructure stack behind it.

You also have other services in the data center, such as backup and virtualization. Your storage should be able to seamlessly integrate with your backup and other tools without having to jump through a bunch of hoops. And, finally, you need a storage system that includes a comprehensive API so that you can fully integrate storage operations into your workflow. For example, you can integrate your storage operations into your existing VMware, Microsoft, Commvault, and OpenStack environments. Or, use the storage systems' REST API, SMI-S, SNMP, Python, or PowerShell capabilities to integrate into just about any tool you can imagine.

Only with a purpose-built all-flash architecture with simple and extensible management tools can you accomplish all of this.

Chapter 7

Ten Reasons You Need All-Flash Storage

.....

In This Chapter

- ▶ Looking at the benefits of all-flash storage
 - ▶ Ending the forklift upgrade
 - ▶ Reducing your storage footprint
-

More than ever, now is the time to make the jump into an all-flash data center. By doing so, you can simplify your data center, which reap rewards throughout the business. You also have all of the performance you can handle. If you're not quite ready to go all-in, you can get easily started with one workload and expand as it makes sense. You're well on your way to taking a bigger leap into building a modern all-flash data center.

This chapter gives you ten (okay, seven) reasons why you need all-flash storage. Of course, these facts are contingent on buying the *right* all-flash solution. Not all are created equal, so do your homework.

Flash Is Fast . . . Really Fast

IOPS? Flash has them in abundance. Throughput? As long as you buy a storage array with controllers that were designed for flash and doesn't use retrofitted disk controllers, throughput challenges will be a thing of the past. Latency? Because IOPS and throughput challenges are effectively eliminated, storage-induced latency is a thing of the past.

Your applications will be more responsive. You're able to consolidate more workloads onto fewer storage arrays because you don't need to worry about performance. All of this means that you will save money.

Flash Reduces Power and Cooling

Less hardware equals less electrical usage. Fewer spinning disks equals less electrical usage. Fewer spinning disks also equates to less generated heat because there are fewer moving parts. As such, flash is not only more reliable but also moving to an all-flash data center can have a major positive impact on your power and cooling bill.

Flash Increases Usable Capacity

For a long time, flash disks didn't hold as much data as spinning disks. Those days are long gone, even when you compare raw capacities between the two types of media. However, with all-flash, you get additional benefits thanks to data reduction. While also available on disk-based systems, it's generally not nearly as effective as reduction in all-flash environments. The result is massively increased efficiency and effective capacity, which helps to drive down the amount of hardware you need and also drives down your effective cost per gigabyte of storage capacity.

Flash Provides a Comprehensive API

Automation is a big deal for companies looking to streamline their IT activities and operations. A robust API that can enable scripting of all routine storage functions can be a huge boost to these efforts. Fortunately for you, an all-flash storage array exists that includes such an API.

Consider the possibilities. Virtualization administrators live in the world of the vSphere web client, so using a storage system that can be managed directly from the web client with no need to leave the application is a huge time saver. With PowerShell and OpenStack integration, you can automate workflows in Microsoft and OpenStack environments, respectively.

By choosing an array with support for REST-based API calls, SNMP, and a CLI, you can integrate the system into just about anything under the sun.

You Can Converge All the Things

You can't run a data center on storage alone. You need other things too, like servers, and networking. With an all-flash storage system that also participates with best of breed infrastructure companies in creating a validated converged infrastructure, you can take the worry out of convergence and bring extreme simplicity and performance to the entire data center stack.

All-Flash Architecture Effectively Ends the Forklift Upgrade

Look out over a period of a decade. How many times did you buy the same terabyte over and over? Every time you rip and replace legacy storage, you're re-buying what you already have. And, beyond the crazy economics this entails, think about the disruption to the business and the opportunity cost with regard to IT staff time.

Consider an all-flash data center with ongoing guarantees around buying storage once and easy expansion and your days of having to bring a forklift to the data center are over. Now, you can simply upgrade what you already have to bring new capabilities and more capacity to bear.

All-Flash Reduces Your Storage Footprint

With all-flash, you don't need to buy extra hardware just to get the performance you need. As you reduce the number of devices you have, you also reduce the number of cables that you need to run in order to tether everything together, further simplifying the data center. Moreover, fewer cables translates to less air movement restriction, which can also reduce cooling costs. All-flash also increases the density of storage arrays, enabling you to reduce entire racks of equipment to handfuls of rack units.

The end result is fewer rack units dedicated to storage. That directly reduces physical data center costs.

Storage is undergoing a fundamental transformation

Don't get plagued with complexity and waste in your storage solutions — look to cloud IT for answers. Choose to adopt smart storage solutions, such as all-flash systems, that provide companies with cloud-like productivity, efficiency, and economics.

- *Discover simplicity — realize the positive impacts on your business*
- *Solve your virtualization and consolidation challenges — implement flash storage and save money*
- *Eliminate storage-induced budget spikes — find a storage solution that provides variable length deduplication*
- *Find the key ingredients for enterprise cloud — achieve storage success*



Open the book and find:

- How to contain your storage costs
- The future of storage consumption and cost
- Ways to achieve scalability and agility
- The top reasons you need all-flash storage

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

WILEY END USER LICENSE AGREEMENT

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