

Brought to you by:



Modern All-Flash Data Centers

for
dummies[®]
A Wiley Brand

Enable hybrid
cloud environments



Discover the benefits of
storage simplicity



End forklift upgrades
and reduce footprint



Calvin Nieh
Mike Kieran
with Lawrence C. Miller

2nd Pure Storage Edition

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-first approach, drives business and IT transformation with a Modern Data Experience that's effortless, efficient, and evergreen. Pure Storage offers two flagship products: FlashArray//X, optimized for structured workloads, and FlashBlade, ideal for unstructured data. With Pure's industry-leading Medallia-certified NPS score of 82, Pure customers are some of the most satisfied in the world and include organizations of all sizes across an ever-expanding range of industries.



Modern All-Flash Data Centers

2nd Pure Storage Edition

by Calvin Nieh and Mike Kieran
with Lawrence C. Miller

for
dummies[®]
A Wiley Brand

Modern All-Flash Data Centers For Dummies®, 2nd Pure Storage Edition

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

Copyright © 2021 by John Wiley & Sons, Inc.

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 and the Pure Storage logo are registered trademarks of Pure Storage. 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: 978-1-119-73555-7 (pbk); ISBN: 978-1-119-73556-4 (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 Manager:

Carrie Burchfield-Leighton

Sr. Managing Editor: Rev Mengle

Acquisitions Editor: Steve Hayes

Production Editor: Siddique Shaik

Business Development

Representative: Molly Daugherty

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	6
How data reduction increases available capacity	6
Facilities improvements: power and cooling	7
Data center storage footprint and cabling	8
Eliminating Storage Software Licensing Headaches and Costs	8
Introducing Modern All-Flash Storage	9
CHAPTER 2: Attaining Simplicity and Manageability	11
Why Simplicity is Imperative: Scaling the Business	11
How Does Storage Simplicity Benefit an Organization?	12
Reducing procurement complexity	12
Repurposing IT resources	13
Reducing the data center maintenance burden	13
Tuning Flash Storage for Application Workloads	14
Enabling All-Flash Storage with Different Types of Flash	15
CHAPTER 3: Virtualization, Containers, and Consolidation	17
Identifying Storage Challenges in a Cloud-Native World	18
Solving Cloud Native Storage Challenges	19
Reduced latency	19
Liberating server resources to improve VM density	20
Deduplication	21
Management integration with the hypervisor	21
CHAPTER 4: Modeling Future Storage Consumption and Costs	23
Understand Long-Term Storage Economics	23
Discover the Benefits of Data Reduction	26
Delivering Consistent Performance	27

CHAPTER 5:	Achieving Scalability and Agility	29
	Using Simplicity to Improve 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 Consistent Application Performance	32
	Leveraging APIs for Automation.....	33
	Supporting Data Growth and New Technology Innovations	34
CHAPTER 6:	Addressing the Cloud	35
	Looking at Flash and Cloud as Key Disruptors	35
	Bridging On-Premises and Cloud Infrastructure.....	36
	Looking to Modern All-Flash Storage as the Foundation for Enterprise Cloud Strategies	37
	Simple scaling.....	37
	Consistent performance	38
	Always-on.....	38
	Easy management from anywhere.....	38
	OpEx cost models	39
	Integration capabilities.....	39
CHAPTER 7:	Ten Reasons You Need Modern All-Flash Storage	41
	Gives You Speed	41
	Accelerates All Your Workloads	42
	Consolidates Database Workloads	42
	Increases Effective Capacity.....	42
	Enables Automation and Orchestration.....	43
	Converges Your Infrastructure	43
	Reduces Power and Cooling Costs.....	43
	Reduces Your Storage Footprint	44
	Ends the Forklift Upgrade	44
	Delivers the Modern Data Experience.....	44

Introduction

The modern data experience enables simple, seamless, and sustainable data storage as a service. By helping you use more of your data while reducing the complexity and expense of managing the infrastructure behind it, modern all-flash storage enables the modern data experience and empowers the data-driven enterprise to extract maximum value from its data.

About This Book

This book consists of seven chapters that explore how

- » Traditional storage costs add up and how modern all-flash storage reduces total cost of ownership (Chapter 1)
- » Storage simplicity and manageability nirvana can be achieved with modern all-flash storage (Chapter 2)
- » Virtualization, containerization, and consolidation are supported by modern all-flash storage (Chapter 3)
- » Storage consumption and costs are optimized with modern all-flash storage (Chapter 4)
- » Cloud scalability and agility are achieved with modern all-flash storage (Chapter 5)
- » Hybrid cloud environments are enabled with modern all-flash storage (Chapter 6)
- » Modern all-flash storage delivers real business and technical benefits (Chapter 7)

Each chapter is written to stand on its own, so if you see a topic that piques your interest, feel free to jump ahead to that chapter.

Foolish Assumptions

It's been said that most assumptions outlive their usefulness, but we assume a few things nonetheless! Mainly, we assume that you have some understanding of data storage technologies and concepts. Or perhaps you work in a storage, virtualization, or data center architecture role. If these assumptions describe you, then this is the book for you. If none of these assumptions describe

you, you should keep reading anyway. This book gives you quite a bit of useful information about the modern data experience.

Icons Used in This Book

Throughout this book, we occasionally use special icons to call attention to important information. Here's what to expect:



REMEMBER

This icon points out important information you should commit to your noggin — along with anniversaries and birthdays.



TECHNICAL
STUFF

This icon explains the jargon beneath the jargon and is the stuff legends — well, legendary nerds — are made of. You can skip this info if you aren't interested.



TIP

Tips are appreciated but never expected — and we sure hope you appreciate these useful nuggets of information.



WARNING

These alerts point out the stuff your mother warned you about (well, probably not), but they could help you avoid a mistake.

Beyond the Book

We can only cover so much in this short book, so if you want to find out more, check out www.purestorage.com/flash and these other helpful resources:

- » **Flash Storage Buyer's Guide:** <https://bit.ly/FlashBuyersGuide>
- » **IDC white paper:** <http://bit.ly/IDCEvergreen>
- » **Forrester Consulting study:** <https://bit.ly/ForresterTEIPure>
- » **Cloud Buyer's Guide:** <http://bit.ly/CloudBuyersGuide>
- » **Real World Stories of IT Admin Heroes:** <https://bit.ly/realworlditheroes>
- » **NVME over RoCE Storage Fabrics for NOOBs:** <http://bit.ly/NVMENoobs>

IN THIS CHAPTER

- » Adding up the TCO 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
- » Learning about modern all-flash storage

Chapter 1

Containing Storage Costs

Data center infrastructure has long been seen as an unavoidable cost of doing business. Companies are always looking for ways to manage these expenses while still delivering the information resources their people need to innovate and create value. For data storage specifically, the technology landscape has changed rapidly of late, giving IT leaders a great opportunity to reassess the cost effectiveness of their current systems.

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 a modern all-flash storage solution can solve it?

Calculating the TCO of Traditional Storage

Clearly defining the problem is an important first step in transforming your data center for greater efficiency. In terms of storage, you need to understand your fully loaded storage costs

associated with traditional or legacy storage infrastructure. The TCO for storage goes far beyond the initial purchase price and ongoing maintenance that you pay. Quite a few fixed and variable costs should be considered:

» Capital expenditures (CapEx)

- **Initial purchase:** 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 component of storage procurement.
- **Replacements and upgrades:** Everything in your data center is on a replacement cycle. The time to decommission those expensive storage arrays will come. Most storage vendors recommend a full replacement or a major upgrade of your equipment in order to stay current with the latest and greatest technology.
- **Additional capacity:** Ensuring sufficient capacity for data growth is a never-ending challenge for storage administrators. This challenge is further complicated by the fact that, for performance reasons, many traditional storage arrays don't allow you to use anywhere near the available drive space within a storage array. As a result, many companies inevitably overspend on expensive drive space and additional array shelves to ensure they always have enough (that is, too much) storage space.

» 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.
- **Power and cooling:** Many companies use a combination of traditional hard disk drive (HDD) arrays and hybrid flash storage arrays that consist primarily of HDDs in their data center environments. With so many moving parts, these "spinning" HDD arrays consume a ton of wattage and generate a lot of heat. 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 means that you're faced with greater 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 software licensing costs. For example, do you want data deduplication? What about cloning? Replication? Encryption? Even with all-flash storage systems, some vendors charge a lot for these features. In fact, licensing costs can add a significant amount to overall storage costs.
- **Floor space:** The traditional HDD storage systems that many organizations still use today simply can't meet modern performance requirements. To increase performance, storage administrators need to add more disks, which results in data center racks that are full of whirring disks, spinning away, and taking up valuable — and expensive — rack space.
- **Complexity:** Many traditional storage systems are truly complex. The more complex they are, the harder they are to manage. The result is excessive IT personnel costs for rudimentary storage management. Even worse, this complexity can lead to slowdowns for new business application deployments, increasing the costs for those projects as well. Complexity is an expensive characteristic in any IT environment. It may be difficult to quantify, but it has significant cost implications.



WARNING

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 perils become more apparent, and the true TCO for traditional storage becomes painfully evident.



TIP

Some companies prefer a cloud consumption model for data storage and other IT resources, which enables them to shift expenses from upfront CapEx to more flexible “pay-as-you-go” OpEx. They pay only for what they use, when they use it, and can scale on demand as their business needs grow.



REMEMBER

The TCO of traditional storage isn't limited to on-premises data centers. Many companies are paying for traditional storage arrays and devices in remote or branch locations at the edge of their networks, as well as in public cloud environments.

Discovering the All-Flash Ingredients that Reduce TCO

In the early days, data center flash was considerably more expensive than “spinning rust.” Those days are long gone. Modern all-flash storage solutions are *far* more affordable than the original Serial Attached Small Computer System Interface (SAS) hybrid flash arrays from more than a decade ago. Modern all-flash storage can help organizations reduce both their direct and indirect storage costs.



TECHNICAL
STUFF

Older hybrid flash arrays use a combination of SAS and Serial Advanced Technology Attachment (SATA) host controller interfaces to connect to solid-state drives (SSDs) and HDDs, respectively. Modern all-flash storage uses Non-Volatile Memory Express (NVMe) host controller interfaces which provide higher efficiency, better performance, and greater performance-to-density ratios than SAS and SATA. Not all storage vendors offer a modern all-flash storage solution built on NVMe technology, so this is an important factor to consider when making a purchase decision.

How data reduction increases available capacity

Even as intelligent data services become a key consideration when choosing storage systems, capacity remains a fundamental concern. Fortunately, every dollar in your budget buys you a lot more storage capacity than it did a few years ago, thanks primarily to vast improvements in data reduction technologies.



TECHNICAL
STUFF

Data reduction is an umbrella term that refers to features that can help organizations store more data in their existing storage. You can effectively reduce the amount of data that you have to store and achieve greater storage efficiency using different methods, such as data deduplication and compression (we cover these items more in Chapters 3, 4, and 6).



REMEMBER

With data reduction, you get a multiplier effect that, in essence, increases your actual storage capacity. Data reduction technologies are specified as a ratio of how much the original data is reduced. This ratio — represented as effective reduction:1 — tells

you how effective the array's reduction capabilities are based on your workloads.

For example, a “10:1” data reduction means that you're actually able to store the equivalent of ten times as much information on every byte of storage in your storage array, effectively increasing its capacity by 1,000 percent. Although data reduction is often an optional feature on traditional storage systems, modern all-flash storage integrates reduction technology right into the array, making it a core part of the array's operating system.

Now, think about this from an economics perspective. For every dollar you're spending on *raw* storage capacity, you're actually getting \$10 worth of *effective* storage capacity. That's just one way that modern all-flash storage can begin to help you tame the beast of storage economics.

Facilities improvements: power and cooling

Facilities costs related to power and cooling are often a significant part of an organization's infrastructure expense. All-flash storage reduces these costs, in part due to some basic physics.

Physics? Yes, physics. Here's why: All those moving parts in a traditional hard disk drive 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 in flight. All-flash storage media has no moving parts. While arrays will still have fans to provide cooling, the flash drives 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.



REMEMBER

With modern all-flash storage, achieving data reduction rates of 10:1 means that you need *much less* hardware than in the past. Data centers using traditional spinning hard disk drives typically didn't have data reduction capabilities by default, so you only had raw disk space to work with. Enabling data reduction by default across an array means 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

Data center rack space can be really expensive. While virtualization has helped to reduce the amount of physical equipment needed, equipment space needs grow over time as new workloads are added.

In the past, storage has been a key contributor to these costs, and it's still a major factor for companies that continue to use inefficient old HDD or hybrid systems. It's also a factor for companies with first-generation flash products that don't use advanced data reduction capabilities because

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

Many organizations have found that flash storage systems effectively eliminate their performance concerns. And on the capacity front, flash disks are increasing in size all the time. Between capacity increases and data reduction features, you may find you don't need as much storage hardware.



TIP

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 translate 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 may 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 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.



REMEMBER

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.

Introducing Modern All-Flash Storage

For too long, you've worried about your storage. You've been uncertain about its capability to keep up with your workload demands, concerned that the manufacturer would suddenly stop supporting your arrays, and unsure whether your budget would cover the software upgrades needed to keep up with evolving business requirements. It's time to change the paradigm.



REMEMBER

With modern all-flash storage, those worries go away as long as the storage purchase is coupled with an all-inclusive array 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.

Modern all-flash storage enables the four pillars of the modern data experience:

- » **Fast matters.** Fast data is agile data. Modern all-flash storage redefines fast with the lowest latency and highest data throughput.

- » **Cloud is everywhere.** Modern all-flash storage bridges the enterprise data center with the cloud and provides industrial strength block storage to seamlessly run existing apps in the cloud.
- » **Simple is smart.** Modern all-flash storage provides intuitive storage management, even as technology environments and businesses grow ever more complex.
- » **Subscribe to innovation.** Modern all-flash storage enables smooth technology evolution with automatic hardware and software upgrades that let you keep pace with technology evolution without downtime or performance impact.

IN THIS CHAPTER

- » Finding out why simplicity is important in storage
- » Discovering the benefits of storage simplicity
- » Optimizing flash storage for your application workloads
- » Leveraging new flash technologies in modern all-flash storage

Chapter 2

Attaining Simplicity and Manageability

Information Technology (IT) can be complex — most folks know that. In fact, in some IT shops, complexity is not only accepted, but also it's glorified a bit. But in modern IT environments, that attitude is as outdated as punch cards. Today, software is eating the world. Modern enterprise applications demand high performance and non-stop reliability, and complexity makes them impossible to provision and manage efficiently.

In this chapter, you discover why simplicity has become imperative across all aspects of the data center and how simplicity in your storage infrastructure can have positive impacts across many different aspects of the business.

Why Simplicity is Imperative: Scaling the Business

There was a day when deploying data center products that had complicated administrative interfaces may 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 and ease of use
- »» Streamlined administrative experience
- »» Automation
- »» Simple procurement process

Organizations of all sizes are proactively seeking ways to make IT infrastructure simpler to manage. Business and IT leaders want to ensure that budgets and resources are invested in areas that contribute to innovation and digital transformation. And to increase data center efficiency, new technologies must be easy to configure and deploy.

Traditionally, storage has been among the most expensive and complex resources managed in the data center, so it's no surprise that modern all-flash storage arrays have become the new standard. All-flash arrays have eliminated storage performance challenges by offering breakthroughs in speed, scalability, simplicity, and agility — not to mention being easier to procure, manage, and maintain than traditional disk-based storage systems.



REMEMBER

Simplicity enables an organization to more quickly scale its business.

How Does Storage Simplicity Benefit an Organization?

Traditionally, storage is one of the most complex infrastructure components in a data center; therefore, making storage management simpler can enable a number of significant benefits for your company.

Reducing procurement complexity

Simplicity reduces the amount of time you spend procuring storage. Here's proof: Were you to compare a bill of materials for a storage purchase from a modern all-flash storage vendor and a traditional storage vendor (even a traditional storage vendor with all-flash options), the modern all-flash storage vendor will

probably show just two to three items to consider. That's simplicity at work. On the other hand, traditional storage vendors may actually have several pages of options, products, and services. That's complexity ad nauseum. So, right from the start, a simple modern all-flash storage solution begins to save you money in terms of procurement complexity.

Repurposing IT resources

Everyone needs 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 devote 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 will 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.



WARNING

Over the life of a traditional storage array, maintenance costs can become a real albatross. Hardware maintenance works a lot like life insurance. When you first buy a traditional storage array, your maintenance costs may be relatively low. Over time, though, vendors raise the cost of maintenance as the risk of a catastrophic hardware failure in your aging traditional storage array increases. Eventually, the cost of maintenance — along with gradually deteriorating performance and technology obsolescence — makes outright replacement more palatable, albeit still expensive. In other words, you have to rebuy your storage every few years and then pay an increasing amount annually to support and maintain it, despite diminishing returns on your investment in terms of performance and innovation.



TIP

Evergreen Storage is an innovative subscription program from Pure Storage that eliminates expensive storage technology refreshes and ever-increasing maintenance costs. The Evergreen Storage experience continuously delivers Pure Storage innovations in software and hardware to you — without any disruption, downtime, or data migrations. It's like getting the benefits of software-as-a-service (SaaS) and the cloud in your on-premises storage deployment. Discover more about Evergreen at www.purestorage.com/products/evergreen.html.

Tuning Flash Storage for Application Workloads

If you're buying an all-flash array, it's likely that you want to maximize your return on investment (ROI) 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 applications 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 re-imported to conform to the required input/output (I/O) size; this process is a common recommendation with online transactional processing (OLTP) databases, virtual desktop infrastructure (VDI), and virtual machine workloads.

Worse still, 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 because 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 modern all-flash technology will always outperform a retrofitted system.



REMEMBER

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. Some flash arrays also impose severe restrictions on the data layout for your applications, both during initial configuration and when being re-configured to accommodate growth.

Enabling All-Flash Storage with Different Types of Flash

Not all application workloads and storage use cases require the fastest possible performance. For example, archiving, disaster recovery, development, and testing are typically considered “tier 2” workloads, with their data allocated to lower cost hard disk drives (HDDs) — in particular, hybrid storage arrays with slow-speed and relatively unreliable Serial Advanced Technology Attachment (SATA) disk drives.

Single-level cell (SLC) is the fastest, most reliable, most accurate and, not surprisingly, most expensive type of flash storage technology commonly available today. The most widely used type of flash storage technology is multi-level cell (MLC). In the early days of data center flash storage, MLC brought a steep drop in price because it was capable of storing multiple (two) bits per flash cell. However, compared to traditional HDDs, flash storage was still cost prohibitive for most tier 2 storage use cases. Therefore, hybrid storage arrays remained firmly entrenched in the enterprise data center.



In flash technology, a cell is a floating-gate transistor that has two possible states (or voltage levels) representing one bit of data (1 or 0). The more cells there are in a particular flash storage technology, the higher the potential data density (capacity), which means lower cost but also means shorter life (fewer program/erase, or P/E cycles) and a higher bit error rate.

The next evolution in flash storage technology was triple-level cell (TLC) flash. TLC flash technology makes all-flash storage more cost-effective, and therefore more feasible, for tier 1 and tier 2 storage requirements. TLC is particularly fortuitous given the exponential growth of enterprise storage requirements due the tendency to never delete anything, and the trend toward implementing Agile methodologies and a DevOps culture.

Quad-level cell (QLC) delivers the lowest cost/GB across flash technologies and makes it extremely viable for all-flash storage to replace hybrid storage arrays. However, most traditional storage vendors are not able to effectively utilize QLC because their legacy architectures can't directly communicate with NAND flash in an efficient manner.



TECHNICAL
STUFF

NAND flash gets its name from the fact that it connects floating-gate transistors in a manner that resembles a NOT-AND (NAND) logic gate.



WARNING

If a storage vendor tells you it doesn't use QLC because it isn't reliable enough, it's too slow, or the cost difference between QLC and TLC isn't significant, look for a different storage vendor. The real reason is likely that its legacy architectures don't support this innovative new technology.

Table 2-1 summarizes the differences between SLC, MLC, TLC, and QLC flash technologies.

TABLE 2-1 Comparison of Flash Technologies

	SLC	MLC	TLC	QLC
Performance	Highest	High	Medium	Slowest
Flash cell (P/E) cycles	50,000 to 100,000	1,000 to 10,000	500 to 1,000	150 to 1,000
Bit error rate	Low	Medium	High	Highest
Price per GB	High	Medium	Low	Lowest

Storage Class Memory (SCM) is a new tier of memory/storage that sits between dynamic random-access memory (DRAM) and NAND flash in terms of performance, endurance, and cost. Unlike DRAM, SCM is persistent in nature and retains data written to it across power cycles. Compared to NAND flash, SCM is orders of magnitude faster for both read and write operations. Unlike NAND, SCM is significantly more resistant to data re-writes and has much higher endurance properties. It delivers these benefits over NAND flash at a much lower cost/GB as compared to DRAM.

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

Chapter **3**

Virtualization, Containers, and Consolidation

The cloud has ushered in a new era in which traditional data center resources are increasingly virtualized. Automation and orchestration, Representational State Transfer (REST) application programming interfaces (APIs), and other modern innovations have given rise to cloud-native technologies such as containers and microservices. At the same time, organizations are rapidly consolidating their on-premises infrastructure with virtualization and software-defined architectures, while moving more of their traditional on-premises data center resources — including compute and storage — to multicloud and hybrid cloud environments.

In this chapter, you discover how modern all-flash storage supports these important trends and technologies.

Identifying Storage Challenges in a Cloud-Native World

Virtualization is a key enabling technology in the cloud and has also played a pivotal role in on-premises data centers. Organizations virtualize infrastructure for many reasons, including to

- » **Reduce costs:** Virtualization can be used to consolidate multiple workloads that were formerly housed on individual servers and helps reduce the amount of hardware needed.
- » **Increase utilization:** Virtualization enables underutilized compute and storage resources to be consolidated and managed as shared resources instead of operating in disparate silos with wasted capacity.
- » **Simplify management:** With less physical hardware comes less overall management complexity.

Unfortunately, not everything is ideal when it comes to virtualization and the cloud. Some key storage challenges must be overcome for the full benefits of virtualization, the cloud, and cloud-native technologies to be realized.

The old days of virtualization storage used all-flash arrays that relied on Serial Attached Small Computer System Interfaces (SAS) controllers, which constrained their performance (the really old days of hybrid arrays containing spinning hard disk drives, are now considered ancient history).

SAS evolved from Serial Advanced Technology Attachment (SATA) controllers as storage vendors began adding solid-state drives (SSDs) alongside spinning hard drives in hybrid arrays. SAS was a definite upgrade from SATA and needed to take advantage of the significant performance improvements of flash storage over hard disks.

For example, the SATA III interface is limited to running at 6 gigabits per second (Gbps) with up to 600 megabits per second (Mbps) of throughput, largely due to its half-duplex operation — it can't perform read and write operations simultaneously. A SAS interface, by comparison, runs at 12 Gbps and delivers up to 8 Gbps of bandwidth throughput. In addition to performance improvements, SAS also enables full-duplex operation (simultaneous

reads and writes), built-in error checking for end-to-end data integrity, and the ability to connect as many as 128 SSDs.

SAS made it possible for storage vendors to leverage both spinning HDDs and flash drives in their hybrid storage arrays, but like most hybrids, it sacrifices some of the benefits of the newer technology for the sake of backwards compatibility with the older technology (think about your old boom box with a CD player *and* a cassette player, or your hybrid car that isn't quite a Tesla). The solution to this problem is Non-Volatile Memory Express (NVMe).

NVMe may sound familiar. Your computer's memory is "non-volatile memory" and it plugs directly into the motherboard. NVMe works in much the same way. It eliminates the bottleneck — the SAS controller — between your storage and the virtual servers running your workloads. NVMe provides throughput of 1 Gbps per lane (16 Gbps in a 16-lane configuration).

Solving Cloud Native Storage 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).

For instance, provisioning storage systems has traditionally been a slow and complicated process, which simply doesn't fit with the concept of containerized deployments. Modern cloud-native storage systems overcome this constraint with robust tools for automated provisioning. Similarly, containerized environments must be able to rapidly scale with automated multi-array provisioning.

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 — input/output operations per second (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.



REMEMBER

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 was originally designed for disk and has been retrofitted to support flash, 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.



REMEMBER

Storage Class Memory (SCM) is a new option for further reducing latency in some all-flash storage systems. New technology such as Intel Optane SCM chips make it possible to get latency as low as 150 microseconds (μs) for your organization's most critical workloads.

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. Modern storage systems provide robust REST APIs that empower users to customize and automate virtually every aspect of their infrastructures. And the rapid growth of private cloud makes companies focus more on integration with virtualization solutions from VMware, Microsoft, and others.

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 and container orchestration management tools makes it easier to adapt to changing performance demands of new workloads.

IN THIS CHAPTER

- » Discovering how to avoid blowing your storage budget
- » Increasing storage capacity by reducing data
- » Ensuring consistent storage performance while reducing costs

Chapter 4

Modeling Future Storage Consumption and Costs

Deploying a modern all-flash storage solution can give your company much greater insight and control over its data storage costs. 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

In Chapter 1, I discuss 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 ten-year period, the numbers may be downright frightening. During that time, you may have bought a storage array and then replaced it at least once. You also paid for licensing, maintenance, support, and additional capacity. If you had to migrate data from an older system to a newer one, you may have paid for professional services, and there was likely also a business cost associated with downtime during the migration.

All this talk of replacement and cost can be a bit scary, but there’s a solution. Modern all-flash storage can drive major savings in the storage equation and make your matrix look more like what’s shown in Figure 4-1.

<i>CapEx</i>	<i>OpEx</i>
Purchase price	Maintenance: Predictable and fair Power and cooling: Massively reduced Floor space: More storage density equates to less floor space, which equates to lower costs 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 No capacity surprises: Guaranteed effective capacity, based on your initial storage workloads; protects your storage investment

FIGURE 4-1: A storage environment built on evergreen storage principles.

In the world of IT, major CapEx purchases need to be budgeted and approved every time you replace a resource as it nears the end of its useful life — around every three to five years. Unfortunately, in the world of finance, fixed assets typically depreciate on a five- or seven-year schedule. Therefore, the definition of useful life differs between IT and finance, and never the twain shall meet.

Even worse, major CapEx purchases are also necessary when things go wrong — and these can be your most costly purchases. If a storage array fails or you don’t have sufficient storage space due to poor capacity planning, you need to act quickly. You don’t necessarily have time to shop around for the best technology or negotiate the best deals — you’re essentially panic buying. Unbudgeted contingency spending has a high direct cost, and a

potentially higher indirect cost because it may cause budget overruns and delays in other projects (such as application development) that are reliant on storage capacity, and because of the opportunity cost (that is, you have to forgo or delay some other project that may have had a high return on investment) incurred due to the unplanned spending.



REMEMBER

When you replace your storage, you're also subjecting the company to some level of business disruption. You need to schedule downtime so you can move your application and data to the new storage platform — and that downtime has a real business cost.

CFOs *hate* large CapEx budget increases like the one shown in Figure 4-2.




FIGURE 4-2: An exaggerated CapEx budget increase . . . for effect.

With a modern all-flash storage solution, you can eradicate major CapEx budget increases and contingency spending on storage, once and for all. After your initial storage purchase, you can keep it for ten years or more and reliably predict your costs over the lifetime of the solution.



TIP

With the Pure Storage Evergreen Storage guarantee, you never have to rebuy a terabyte (TB) you already own. Evergreen Storage provides the benefits of software as a service (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, customers can expect to avoid two to three upgrades over a ten-year life cycle.

Discover the Benefits of Data Reduction

Data reduction technology constantly evolves. So, it should come as no surprise that, as you review your all-flash storage options, you find that different arrays provide different results depending on the data reduction technologies they use and how they implement them. In a modern all-flash storage solution, these technologies should include

- » **Pattern removal:** Pattern removal identifies and removes repetitive binary patterns, including zeroes. In addition to capacity savings, pattern removal reduces the volume of data to be processed by the dedupe scanner and compression engine.
- » **512-byte aligned variable dedupe:** High-performance, inline deduplication operates on a 512-byte aligned, variable block size range of 4 to 32 kilobytes. Only unique blocks of data are saved on flash — removing even the duplicates that fixed-block architectures miss. Best of all, these savings are delivered without requiring any tuning.



TIP

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

- » **Inline compression:** Inline compression reduces data to use less capacity than the original format. Append-only write layout and variable addressing optimize compression savings by removing the wasted space that fixed-block architectures introduce. Combined with deep reduction, compression delivers two to four times the data reduction, and is the primary form of data reduction for databases.
- » **Deep reduction:** Beyond inline compression, additional, heavier-weight compression algorithms are applied post-process that increase the savings on data that was compressed inline. Unlike modern all-flash storage, most other all-flash products don't use multiple compression algorithms and simply miss these savings.



TIP

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 on to provide a quantifiable advantage in data reduction over those with only a single form.

» **Copy reduction:** Copying data on a flash array only involves metadata. Leveraging the data reduction engine, modern all-flash storage provides instant pre-deduplicated copies of data for snapshots, clones, replication, and xcopy commands.



REMEMBER

Inline deduplication is essential for improving reliability in all-flash storage systems by avoiding unnecessary writes and eliminating Program/Erase (P/E) cycles for data that already exists.

Most all-flash vendors claim some level of “effective capacity,” defined as the capacity you’re able to achieve with an array once you take data reduction into consideration. However, like data reduction technologies, not all workloads are created equal and neither are all storage arrays.



REMEMBER

The Right Size Guarantee in Pure Storage’s Evergreen Storage program provides an effective capacity guarantee (not just an efficiency ratio like other storage vendors) based on the kinds of workloads that you’re running. The Right Size Guarantee effectively answers your most important sizing question, “Will I have enough storage from all my data?” with a definitive “Yes” by shifting the sizing risk at first purchase from the customer to Pure.

Delivering Consistent Performance

Capacity and economics are important, but the entire point of storage is to support application workloads. Perhaps the biggest benefit that a modern all-flash storage solution brings to the table is the ability to run application workloads with consistent and reliable performance over its entire life cycle.

From the perspective of your users, consistent performance is, perhaps, the best benefit . . . as long as it’s also fast. How fast? 150 microsecond (μ s) latency in Non-Volatile Memory Express (NVMe), NVMe Over Fabrics (NVMeoF), and Storage Class Memory (SCM) enables consistent, fast performance that is

comparable to direct-attached storage, but with the efficiency of shared storage. This means modern all-flash storage can support the most demanding workloads whether on-premises or in the cloud, including databases, virtualization and containers, virtual desktop infrastructure (VDI), enterprise resource planning (ERP), electronic health records (EHR), and so on.



TECHNICAL
STUFF

A microsecond is one millionth of a second, or 1/1000th of a millisecond.

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.



REMEMBER

As you consider all-flash storage options, you also need to be able to recover from hardware failures. Any failures that you endure shouldn't be debilitating to the business. With a modern all-flash storage solution, if you've suffered a controller loss, for example, you can maintain full performance. The same is true if you have to take a controller temporarily offline for an upgrade. Some flash arrays experience a 50 percent performance degradation during upgrades under load. In such instances, customers are likely to leave that array partially unloaded in order to reduce the impact of upgrades. With a properly designed modern all-flash storage solution, there is never a performance degradation during upgrades, or even failures — and you get full value for your investment.

IN THIS CHAPTER

- » Bringing simplicity and efficiency to storage
- » Ensuring a consistent modern data experience
- » Improving efficiency and agility with automation and orchestration
- » Scaling for growth and leveraging the latest technology

Chapter 5

Achieving Scalability and Agility

IT departments are routinely expected to ramp up service levels without increasing costs. At the same time that the business asks for more services and tighter service level agreements (SLAs), it also expects that IT protects 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 to Improve Scalability and Agility



TIP

Your storage systems provide a great example of how increasing operational simplicity can make your business more agile and scalable. Help your business forge ahead by providing a modern all-flash storage 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 much more than just provide a place for you to store your data. In fact, with the right agile data services, storage can be a critical component of your organization's strategic capabilities.

Consolidating storage, data reduction, and disaster recovery makes your IT operations a lot simpler, which can free up budget that you can use for other business-facing needs.

Adopt converged infrastructure

There's more than just storage in your data center and cloud environments. All these additional resources create more complexity. Making everything work together can often be time-consuming. Moreover, integrating solutions from different vendors can hamper deployment of new business applications. Technical complexity reduces business agility.



TIP

Converged infrastructure — unifying servers, storage, fabric, and a hypervisor — was developed to solve some of these challenges. Unfortunately, many converged infrastructure solutions are based on old, slow, inefficient hard disk drive (HDD) storage or on first-generation flash systems. The resulting architectures are too expensive, perform poorly, or have significant downtime — making them unsuitable for the consolidation of business-critical applications and workloads.

An effective converged infrastructure solution must provide

- » High levels of performance and scalability
- » Simple to operate and doesn't add unnecessary complexity to the storage environment
- » Enterprise-class equipment to enable dependable support for all applications
- » Economically viable and able to help you reduce your overall infrastructure costs while also providing the full benefits of an enterprise storage solution



REMEMBER

By deploying the right converged infrastructure solution, you make it simple to scale your storage environment without having to worry about excessive costs or extended deployment times. 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. Native integration with many other third-party monitoring, analytics, and management tools (such as Cisco Intersight) also reduces time to value and helps improve operational efficiency.

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 data reduction achieved, 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. With information at your fingertips, formulating data-driven decisions is easier.

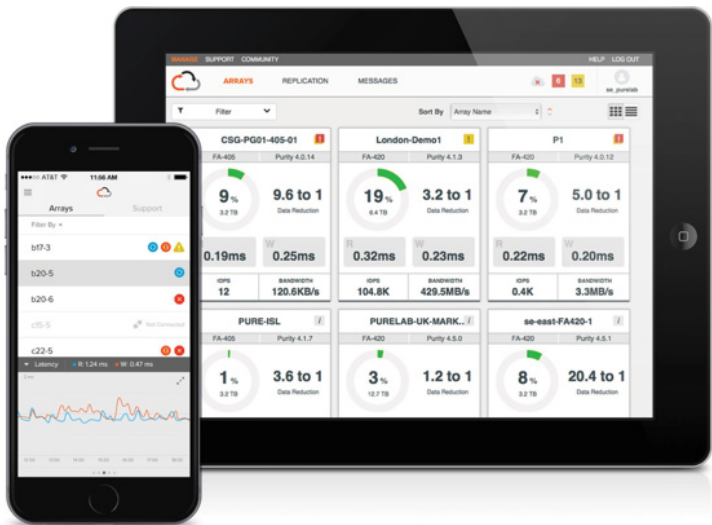


FIGURE 5-1: Pure Storage Pure1 administrative interface provides real-time at-a-glance details.

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 unsupported or near their end of life. These scenarios can extend deployment times and add significant extra costs.



TIP

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. Storage should never be a stumbling block.

Driving Scale with Consistent Application Performance

Consistency is often more important than raw performance. After all, consistency is what drives certain guarantees made to the business, such as service-level agreements (SLAs). Performance consistency comes in two flavors. First, it means that you can guarantee certain levels of input/output operations per second (IOPS), throughput, and latency. However, performance often refers to availability, too. In the case of storage, both definitions of the term are important.

On the performance front, your storage system should provide a steady, consistent experience. For example, performing a software upgrade or replacing a controller should have zero impact on operations. 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's a lot of wasted capacity (and money) to account for a potential hardware failure or occasional upgrade, and it isn't a great recipe for scalability. Only when you're able to guarantee that performance levels will always remain consistent can you

comfortably consolidate applications on less hardware to improve scalability and reduce your costs.



REMEMBER

Having good monitoring ability helps with consistency, but so does a system that can deliver high availability. Look for a system with no single point of failure to ensure resilience and minimal downtime with easy, non-disruptive plug-and-play (no tuning required) upgrades — whether for additional capacity, performance, software features, or even faster media (such as Storage Class Memory, or SCM) — and no data migration required.

Leveraging APIs for Automation

Automation is a key requirement in modern all-flash storage. Representational State Transfer (REST) application programming interfaces (APIs) provide easy integration with your apps and help you automate your storage environment with powerful management and orchestration tools. APIs enable DevOps teams to quickly and easily deploy repeatable “infrastructure as code” in which everything, including compute and storage, is easily programmable.

By having the ability to automate operations, your organization can eliminate many routine tasks, which significantly decreases the number of times that a human has to 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 it.



TIP

Pure Service Orchestrator (PSO) is a software layer that allows fleets of Pure Storage FlashArray and FlashBlade storage systems to be federated and consumed through a simple Storage-as-a-Service API. Given the structural preference that container orchestrators and the cloud-native community have for consuming services and isolating the application layer from the underlying infrastructure, PSO makes using Pure modern all-flash storage a truly differentiated, cloud-native experience.

Supporting Data Growth and New Technology Innovations

As your business grows, so too do its data storage requirements. Scalability is important when evaluating storage solutions, but the main factor usually isn't capacity.

First, you have to consider effective capacity, not raw capacity. *Effective capacity* is what you get after you apply data reduction. For example, a 20 terabyte (TB) array with 10:1 reduction will have an effective capacity of 200TB. Of course, different workloads will yield different reduction results.

Today, most all-flash arrays periodically dial back to the vendor mothership and upload metadata about the overall health of your array and hints at what applications you're running. After analyzing all of this data, storage vendors can begin to provide you with some guidance as to what you should experience with your own application mix.

The ability to quickly leverage new technology innovations is an important key to business agility and competitiveness. In the storage industry, technology innovations such as Non-Volatile Memory Express (NVMe) and NVMe over Fabrics (NVMe-oF) enabled businesses to consolidate siloed, direct-attached storage (DAS), and new media types like quad-level cell (QLC) flash promise to increase capacity density even further. Chapter 2 has more details on QLC flash.



TIP

Pure Storage offers both latency-optimized (//X) and capacity-optimized (//C) modern all-flash storage solutions to address Tier 1 and Tier 2 workloads — and more. With a modern all-flash storage solution, as your business grows and needs more capacity, and as technology continues to advance, you won'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.

IN THIS CHAPTER

- » Understanding how flash and the cloud have changed computing
- » Recognizing key differences between on-premises and cloud infrastructure
- » Identifying enterprise cloud requirements for modern all-flash storage

Chapter 6

Addressing the Cloud

Cloud computing is here to stay, and it's having major impacts on your data center and your IT organization. This chapter explores the cloud in depth and looks at the reasons that modern all-flash storage makes a great foundation for all your storage needs, whether on-premises or in the cloud.

Looking at Flash and Cloud as Key Disruptors

Flash storage and cloud computing are among the most disruptive technologies since the advent of 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, businesses can get all the performance they need by deploying modern all-flash storage in their data centers and in the cloud. Better yet, this performance doesn't require a capacity compromise (check out Chapter 5 for more info).

The cloud — largely enabled by virtualization and containers — is by far the most transformative technology trend of the last decade. It has completely redefined the way businesses compete

today and in the future. As customers have become more accustomed to an on-demand and need-based experience, the flexibility and scalability of the cloud have allowed businesses to realize competitive advantages in a much larger and highly competitive global market.

Infrastructure can live today in all kinds of different environments, including multicloud, hybrid cloud, managed service providers, software-as-a-service (SaaS), platform-as-a-service, infrastructure-as-a-service (IaaS), and others. No longer is the business limited to deploying new applications in an on-premises data center; instead, businesses can leverage cloud and other service providers to stand up new services.



WARNING

Some in IT have approached the cloud with trepidation because it poses a risk to business as usual in IT. Cloud computing requires new skill sets and talent, and introduces different operating models and challenges. Cloud isn't competition; rather, it empowers IT to take a consultative role with the business to help develop a plan for leveraging both on-premises and cloud resources strategically to deliver a competitive advantage.

Bridging On-Premises and Cloud Infrastructure

On-premises and cloud environments have different management, financial, and consumption characteristics. With on-premises infrastructure, there's dedicated hardware with CapEx, depreciation, and manual management processes. In the cloud, resources are virtualized and shared in a multi-tenancy setting, on-demand consumption is OpEx-based, and everything is self-service with extensive automation and orchestration driven by open application programming interfaces (APIs). The environments also have different application and storage architectures. Cloud storage is relatively simple but easily scaled, whereas on-premises environments feature dedicated storage arrays with rich feature sets.

In order to leverage both on-premises and cloud to experience the best of both worlds, companies try to bridge the divide in various ways. Some invest in software-defined infrastructure to abstract

applications and data from their underlying hardware and make them portable. Others try to make their on-premises data centers more cloud-like by deploying technologies such as virtualization, containers, and automation.

To improve agility and innovation, companies increasingly operate in a complex mix of public cloud, private cloud, and on-premises environments. Strategically, leveraging the cloud is no longer an option but an imperative when it comes to digitally transforming your operations and shifting to an OpEx spending model. Businesses want the agility, elasticity, advanced services, and OpEx convenience of the cloud. They're also looking for the same user experience in performance, reliability, availability, and self-service, regardless of where the data lives.

Workload portability is fundamental to a business's overall cloud strategy. As your organization evolves, so too do its technology requirements. Data silos cause fragmentation that limits visibility and collaboration. They also create friction across a broad range of use cases such as data analytics, regulatory compliance, and business continuity. Your organization needs a solution to simplify data portability and to deliver a consistent user experience, regardless of where the data lives. This means having the flexibility to easily move data between on-premises and multi-cloud environments — now and in the future.

Looking to Modern All-Flash Storage as the Foundation for Enterprise Cloud Strategies

With an understanding of the different characteristics of on-premises infrastructure and the cloud, your organization can make your IT systems more flexible and more cost effective.

Simple scaling

A key ingredient for cloud success is the ability to quickly and easily scale the environment, both in terms of capacity and storage efficiency. As business demands change, your company should be able to easily scale your storage to meet the need. This scalability can come from either adding capacity or from enabling

data reduction services, such as deduplication and compression, to increase the overall effective capacity of the environment. A modern all-flash storage solution that is purpose-built to support flash will provide this kind of scalability.

Consistent performance

For storage to deliver consistent performance, it needs to keep working even when a controller fails or is taken offline to perform an upgrade. Just as important, as you scale storage capacity, performance must keep pace.

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

Always-on

Downtime isn't acceptable. Beyond providing consistent 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 in some environments can have business consequences just as severe as an actual outage.



REMEMBER

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. Your data storage systems need to be always on, even when a component is down for servicing, without any impact to user or application performance. That is the kind of architecture that you need for your 24/7 cloud.

Easy management from anywhere

Do you have complete control over your data storage systems when you're away from the office? Can you manage everything from your laptop, phone, or tablet?

Today's end users expect their corporate applications to have the same instant response and slick user interface as the most popular consumer apps. How does your storage vendor measure up?



WARNING

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 held back by your storage vendor. You're losing precious time if you're forced to get to a PC when something goes wrong.



TIP

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 a modern all-flash storage solution.

OpEx cost models

In an on-premises data center, there will always be a combination of CapEx and OpEx expenses that need to be considered. But the data center economic model is changing dramatically, primarily due to budget priorities shifting to the cloud.

With modern all-flash storage, you can begin to provide cloud-like services to individual business units using a fractional consumption, OpEx-focused economic model. Your business can thus gain the internal cost benefits of cloud computing.

Integration capabilities

Storage is just one of many resources in the data center: there's also compute, the hypervisor, and the fabric to think about. Your organization needs a storage system that can integrate neatly and completely into a converged infrastructure stack. Moreover, it needs a storage system built on top of an *intentionally designed* converged infrastructure stack.



REMEMBER

Storage systems should be seamlessly integrated with your backup and other tools without you having to jump through a bunch of hoops. And, you need a storage system that includes a comprehensive API so 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' Representational State Transfer (REST) API, Storage Management Initiative Specification (SMI-S), Simple Network Management Protocol (SNMP), Python, or PowerShell capabilities to integrate into just about any tool you can imagine.

Only with a purpose-built, modern all-flash storage architecture and a software-defined storage solution for cloud data portability with simple and extensible management tools can you accomplish all of this.

IN THIS CHAPTER

- » Improving performance and increasing capacity
- » Supporting DevOps and containerization
- » Ending the forklift upgrade and reducing your storage footprint
- » Realizing the modern data experience vision

Chapter 7

Ten Reasons You Need Modern All-Flash Storage

In this chapter, I highlight the key reasons for putting modern all-flash storage to work for your business.

Gives You Speed

Modern all-flash is fast . . . really fast! Read and write performance in traditional spinning hard disk drives (HDDs) is rated in milliseconds and generally limited to less than 300 read and write input/output operations per second (IOPS). By comparison, flash performance is rated in microseconds and IOPS are measured in the tens of thousands to millions. But to take advantage of all that speed and performance, you need to move up from disk-only or disk-flash hybrid systems to modern all-flash storage.



TIP

Modern all-flash storage solutions effectively eliminate IOPS and throughput limits, so storage-induced latency is a thing of the past. Your applications will be more responsive, and you can consolidate more workloads onto fewer storage arrays because you don't need to worry about performance. All of this means that you can save money.

Accelerates All Your Workloads

Modern all-flash storage systems leverage the latest innovations in flash technology, such as NVMe and storage class memory (SCM), which is nearly as fast as direct-attached storage (DAS). The list of workloads and use cases where modern all-flash storage adds significant value includes databases, private clouds, containers, analytics, data protection, disaster recovery, artificial intelligence (AI), streaming media, augmented reality, virtual reality, gaming, and more. This means the latest generation of cloud-native databases can run on fast, efficient, and shared storage, instead of sitting in isolated silos that are expensive and hard to manage.

Consolidates Database Workloads

Modern all-flash storage solutions can provide latency as low as 250 μ s while still maintaining high throughput. This level of performance means database server CPUs have to wait less time for input/output (I/O) operations to complete, which potentially reduces the number of CPU cycles and associated database licenses required. NVMe all-flash solutions deliver the performance required by mission-critical tier-1 databases to all your databases — including mixed workloads. If NVMe is also used as the data transport protocol (that is NVMe over Fabrics, or NVMe-oF), connectivity between servers and modern all-flash storage solutions is on par with DAS.

Increases Effective Capacity

In the early days, flash storage couldn't match the capacity capabilities of spinning HDDs. Those days are long gone, even when you compare raw capacities between the two types of media. However, with modern all-flash storage, you get additional benefits thanks to data reduction. While data reduction technologies are also available on disk-based systems, they generally aren't nearly as effective as in modern all-flash storage systems. The result is massively increased efficiency and effective capacity, which helps to drive down your effective cost per gigabyte of storage capacity.

Enables Automation and Orchestration

Automation is a big deal for companies looking to streamline their IT operations — that is, all of them. A robust Representational State Transfer (REST) application programming interface (API) in modern all-flash storage solutions enables automation and orchestration of all routine storage functions. This enables IT organizations to re-architect infrastructure as code to support DevOps continuous integration (CI)/continuous delivery (CD) pipelines and containerization.



TIP

Choosing a modern all-flash storage solution with support for REST-based API calls, simple network management protocol (SNMP), Storage Management Initiative Specification (SMI-S) and a command-line interface (CLI) enables you to integrate your storage with just about anything under the sun.

Converges Your Infrastructure

You can't run your applications on storage alone. You need servers and networking too, among other things. With a modern all-flash storage solution that participates in a broad ecosystem of best-of-breed infrastructure companies to create validated converged infrastructure solutions, you can take the worry out of convergence and bring extreme simplicity and performance to the entire infrastructure stack.

Reduces Power and Cooling Costs

A modern all-flash storage solution takes up less physical space in your on-premises data center environments than a traditional storage array filled with HDDs. Those spinning HDDs require more power than flash technology and generate more heat. Which requires more cooling. Which requires more power. Which requires more money. Moving to modern all-flash storage will significantly reduce your on-premises data center power and cooling costs, or you can migrate your on-premises storage arrays to a modern all-flash storage solution in the cloud — and eliminate your on-premises storage-associated power and cooling costs altogether.

Reduces Your Storage Footprint

With a modern all-flash storage solution, you don't need to buy extra hardware just to get the performance your business applications need. Modern all-flash storage increases the density of storage arrays, enabling you to eliminate entire racks of equipment. The end result is fewer rack units dedicated to storage, which directly reduces your physical storage footprint and reduces your overall data center costs.

Ends the Forklift Upgrade

Every time you rip and replace legacy storage, you're re-buying what you already paid for. And, beyond the crazy economics this entails, think about the disruption to the business and the opportunity cost with regard to IT staff time.



TIP

Purchase a modern all-flash storage solution with ongoing guarantees around buying storage once and easy expansion in the future, and your days of forklift upgrades are over.

Delivers the Modern Data Experience

Modern all-flash storage delivered as-a-service allows customers to extract maximum value from their data while reducing complexity and expense — that's the modern data experience. The modern data experience is also

- » **Simple:** API-defined storage services, common management tools, and actionable analytics reduce complexity and accelerate time to value.
- » **Seamless:** Storage services can handle multiple protocols, tiers, and clouds in a single environment.
- » **Sustainable:** Customers buy only what they need, when they need it, and can upgrade to the latest innovation without pain or penalty.



REMEMBER

Pure delivers the modern data experience — simple, seamless, and sustainable — offers all its core products as a service today.



Tap Your Inner Genius With Smart Storage for the Win

**Simple Is Smart.
Storage That Drives Innovation Is Genius.**

Pure's Modern Data Experience™ helps you use more of your data while reducing the complexity and expense of managing the infrastructure behind it.

To learn more, visit
[PureStorage.com](https://www.purestorage.com)



Storage is undergoing a fundamental transformation

In this second edition of *Modern All-Flash Data Centers For Dummies*, you discover what's really behind a modern data experience — and why it matters. With in-depth looks at the costs of traditional storage and the benefits of simplifying with all-flash, find out how to solve architectural data challenges. Take on virtualization, containers, and hybrid cloud environments with storage that's simple, efficient, and cost-effective. End forklift upgrades for good and enable true data agility with modern all-flash storage.

Inside...

- Reduce costs with all-flash arrays
- Explore benefits of storage simplicity
- Solve virtualization challenges
- Optimize storage and data consumption
- Enable hybrid cloud environments
- End forklift upgrades



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

ISBN: 978-1-119-73555-7
Not For Resale



for
dummies®
A Wiley Brand

WILEY END USER LICENSE AGREEMENT

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