

LEARNING MADE EASY

Western Digital Special Edition

# NVMe Storage

for  
**dummies**<sup>®</sup>  
A Wiley Brand



Learn  
about NVMe

Design your next  
generation infrastructure

Make applications  
thrive with NVMe

Brought to  
you by:

**Western Digital.**

## About Western Digital

Data—big data, fast data, personal data—fuels the accelerated pace of innovation and change in the world. This growing diversity of data opens the door for exciting new possibilities that can only be achieved by developing more precise solutions tailored for specific outcomes.

Western Digital® creates environments for data to thrive. Everywhere data is captured, preserved, accessed, and transformed, it's leading the charge to unlock potential. From mobile sensors to personal devices to advanced data centers, Western Digital's industry-leading solutions create environments where the flow of information leads to smarter decisions, breakthrough discoveries, and deeper connections.

As an industry leader of data storage solutions, Western Digital employs a global workforce focused on delivering the possibilities of data. Its data-centric solutions are marketed under the G-Technology™, HGST, SanDisk®, Tegile™, Upthere™, and WD® brands.



# NVMe Storage

Western Digital Special Edition

for  
**dummies**<sup>®</sup>  
A Wiley Brand

# NVMe Storage For Dummies®, Western Digital Special Edition

Published by  
**John Wiley & Sons, Inc.**  
111 River St.  
Hoboken, NJ 07030-5774  
[www.wiley.com](http://www.wiley.com)

Copyright © 2018 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. Western Digital, the Western Digital logo, G-Technology, SanDisk, Tegile, Upthere, and WD are registered trademarks or trademarks of Western Digital Corporation or its affiliates in the U.S. and/or other countries. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

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

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

ISBN: 978-1-119-50971-4 (pbk); ISBN: 978-1-119-50970-7 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

## Publisher's Acknowledgments

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

**Project Editor:** Carrie A. Burchfield

**Editorial Manager:** Rev Mengle

**Acquisitions Editor:** Steve Hayes

**Business Development Representative:**

Karen Hattan

# Introduction

---

**D**ata storage for top-tier applications is transitioning from hard disk drives (HDDs) to solid-state drives (SSDs). SSDs can deliver much higher performance than HDDs, but that performance is often not realized because the standard interfaces to connect processors to storage were optimized for hard disk storage and don't take full advantage of flash storage performance. Non-Volatile Memory Express (NVMe) is a modern interface protocol that was designed at the outset to take full advantage of the capabilities of flash and even more advanced, future solid-state storage technologies. If you're looking for a way to maximize the performance of top-tier applications, NVMe-based storage deserves your careful consideration.

## About This Book

---

This book consists of seven short chapters that provide an overview of the storage landscape and how it's altered by the advent of NVMe technology. In particular, it examines how the NVMe interface reduces latency and enhances flash storage, to increase application performance by delivering data residing internally or across a data center.

# Icons Used in This Book

Throughout this book, special icons call attention to important information. Take note of these.



REMEMBER

This icon points out information that you may want to remember.



TIP

This icon points out helpful suggestions and useful nuggets of information.



WARNING

Danger! This content helps you avoid costly mistakes.

## Beyond the Book

This book can help you discover more about NVMe-based storage, but if you want more info beyond what this book offers, check out these resources:

- » <http://tegile.com>: Discover IntelliFlash, an Enterprise Flash Storage Platform.
- » [pages.tegile.com/the-business-value-of-nvme.html](http://pages.tegile.com/the-business-value-of-nvme.html): Examine the business value of NVMe and the positive business outcomes NVMe enables.

## IN THIS CHAPTER

- » Keeping up with Moore's Law
- » Persisting data
- » Issues with flash storage
- » Finding a solution

# Chapter 1

# Breaking the Storage Bottleneck with NVMe

Legacy interfaces between the processor and the storage device didn't envision performance-oriented storage media such as flash. As flash-based storage devices and systems become prevalent in the data center, interfaces need to be modernized to accommodate new application performance and latency requirements.

# What about Moore's Law?

Back in 1965, Gordon Moore noticed that in each successive year you could double the number of transistors per integrated circuit. Amazingly, this empirical “law” has held for over 50 years. Nowadays, you get approximately twice the computing power every 18 months.

In accordance with Moore's Law, over time, processors have become ever more powerful. So too has memory. This isn't surprising because the same semiconductor technology produces memory and processor chips.

However, hard disk performance hasn't kept up with that of processors and memory. Hard disks can't rotate any faster without corrupting data. This means that hard disk drives are now the weakest link in the chain, and that impacts the overall performance of a computer system.

## Don't Let Data Disappear

One of the main advantages of hard disk drive (HDD) storage is that the data stored is *persistent*. That means that the data is retained even if the HDD is powered down. This is different from computer *random-access memory*, which is *volatile* and doesn't persist data. If you power down a computer, every bit stored in its memory is lost. This is the reason why before you turn a computer off,



you must make sure that everything you want to keep has been saved to persistent storage media.

## Why can't storage persist?

It would be great if there were a technology that featured the persistence of HDDs but was based on the semiconductor technology that followed Moore's Law of advancement. It turns out, there's such a technology, called *flash memory*. Flash memory is produced on a chip the same way that regular random-access memory is fabricated. However, unlike regular random-access memory, it's a persistent storage media.

Since the introduction of flash, the cost per gigabyte has been reducing steadily. However, it's still a lot higher than the cost of HDD storage, and the capacities of flash-based Solid State Drives (SSDs) pale in comparison with HDDs. That said, especially for application environments where high performance is more important than lower costs, flash-based storage is gaining traction.

## All is not rosy in flash land

There is a problem, however, in realizing the performance that flash-based non-volatile storage could ideally deliver. In legacy computer systems, the *connection* between the processor and the persistent storage device was designed to work with HDDs. Because HDDs

were relatively slow, a high-performance interface circuitry wasn't a requirement. Now, with persistent storage based on much faster flash memory, the performance of the old interface has become a major bottleneck.

## **What's the Solution?**

The solution to improve performance is a new interface technology that's designed from the ground up to connect a fast processor to a fast, non-volatile storage device. Such an interface technology has been devised, called Non-Volatile Memory Express (NVMe). It enables the communication between the processor and the storage device to occur with minimum delay.

Storage solutions with NVMe as an interface are just starting to emerge in the market. Currently, they're being used in the most demanding, performance-intensive applications that require real-time response. As time passes, NVMe will be ubiquitous across all business-critical, primary workloads, and HDD storage will focus on providing low-cost, high-capacity solutions.

## IN THIS CHAPTER

- » Understanding how NVMe is better
- » Looking at features and capabilities of NVMe

# Chapter 2

# The Storage Interface

It's almost always easier to upgrade what you have instead of trying something new — even if you're limiting yourself — and especially if you have a lot of time, money, and effort invested. Legacy Small Computer System Interface (SCSI) and Advanced Technology Attachment (ATA) interface designs have gone through multiple revisions to enhance their performance. However, the need to retain backward compatibility with older computers and storage devices limits how much performance can be enhanced.

# How Is NVMe Better?

Non-Volatile Memory Express (NVMe) is a high-performance, scalable host controller interface that's designed to address the needs of solid-state storage for enterprise and client systems. It wasn't intended to maintain backward compatibility for rotating storage. NVMe supports current flash storage and has the flexibility to work with future memory-class persistent storage technologies as well.

## The PCI express standard

The PCI Express (PCIe) standard is a high-speed serial computer bus standard. It improves on previous bus standards by featuring

- » Higher maximum bus throughput
- » Lower Input-Output (IO) pin count
- » Smaller physical size
- » Better performance scaling
- » More detailed error detection and reporting
- » Native hot swapping capability

These advantages lead to more capable, reliable, maintainable, and scalable connections between processors and peripheral devices, including storage.

## **NVMe as a protocol**

In addition to specifying the physical connections and interface characteristics between a processor and storage, NVMe is also optimized for transporting data from solid-state storage to processors across multiple interconnection types. One such interconnect is the Peripheral Component Interconnect Express (PCIe) bus, a very common internal system connection. Another interconnect is over a network, such as Ethernet or Fibre Channel (FC), commonly known as NVMe over Fabrics (NVMeOF). NVMeOF is a network protocol that transports data across the fabric (network) to the processors.

In addition to being optimized for efficient data transfer, NVMe is a flexible protocol that's able to support future solid-state technologies that supersede present day flash.

## **NVMe versus legacy interfaces**

NVMe was designed to take full advantage of flash capabilities and provides significant performance advantage over both Serial ATA (SATA) and Serial Attached SCSI (SAS) interfaces. Data is transferred between the processor and the NVMe controller via as many as 65,535 queues with as many as 65,536 commands per queue. By comparison, SAS and SATA devices have a single queue that supports up to 254 and 32 commands, respectively. These SAS and SATA limits were adequate to support hard disk performance but fall far short of supporting flash performance.

NVMe queues are designed so that IO commands and associated responses operate on the same processor core to take advantage of the parallel processing capabilities of multi-core processors. This enables each application or thread to have its own independent queue, minimizing the performance impact in the system.



TIP

NVMe's fresh-sheet-of-paper design delivers numerous important benefits:

- » Lower IO latency, due to improved drivers, increased number of queues, and increased number of commands in a queue
- » Improved bandwidth, taking advantage of fast PCIe interfaces
- » Lower CPU overhead when handling IO, leaving more cycles for other computations
- » Able to take better advantage of multi-core processors
- » Boost application response times with higher IO per second (IOPS) at very low latency

Benchmarks have shown that NVMe delivers more than twice as many IOPS per Central Processing Unit (CPU) core, as compared with competing SAS and SATA implementations.

# Additional Features of NVMe

The NVMe interface passes memory blocks rather than SCSI commands, which after all is the purpose of a storage interface. This results in reduced storage access latency compared with other interface protocols.

Hot-swapping a 2.5-inch drive form factor storage device is also supported by NVMe. This enables a storage system to ensure application availability and provide on-demand scalability of data storage resources.

Furthermore, recent processors have implemented NVMe directly on the CPU, eliminating the need for either a host bus adapter or separate IO controller. This not only reduces latencies but also improves reliability. For the same power consumption as a SATA subsystem, an NVMe subsystem can deliver more than five times the performance.



REMEMBER

NVMe delivers the full performance potential of flash storage by efficiently supporting more processor cores, lanes per device, IO threads, and IO queues.

NVMe eliminates wasteful SCSI and ATA IO command overhead. Its simplified command processing is enabled by the fact that all commands are the same size and are located at the same position. Overall, NVMe-based systems are less complex, more efficient, more serviceable, and easier to use than legacy systems based on SATA and SAS.





- » Understanding the deployment strategies for your NVMe storage
- » Considering benefits and drawbacks on NVMe storage implementations

## Chapter 3

# Deployment Models for NVMe Storage

**Y**ou can deploy Non-Volatile Memory Express (NVMe)-based storage in an enterprise IT environment in a few different ways. Pros and cons exist with these options, and which one is best depends on several factors.

# Using NVMe Storage in the Server

One way to deploy NVMe storage is in the server. Traditionally, hard disk drives (HDDs) have been connected to a processor by way of an Input-Output (IO) controller within a computer system. As performance demands for business-critical applications grew, internal HDDs couldn't keep up. In fact, many multi-threaded applications were starved for storage IO.

When flash storage devices or Solid State Drives (SSDs) appeared in the market, this enabled computer systems to co-locate a faster, persistent storage medium close to the processor and deliver the performance required for IO intensive applications. However, due to the inefficiencies of legacy storage protocols and the interface connecting them, the full performance of SSDs couldn't be realized. Because SSD performance was substantially better than HDD, the switch was still justified.

However, as the performance of processor and memory devices continued to improve, legacy storage interconnects have become a real bottleneck. NVMe is designed from the ground up to realize the full potential of solid-state storage media, such as flash and even next-generation storage-class memory (SCM) technology. Adding NVMe-capable flash storage to a server, which plugs right into the Peripheral Component Interconnect

Express (PCIe) bus, enables intensive applications to thrive and provides better response times to the end-user.

Many operating systems already provide support for NVMe and NVMe-capable SSDs. Therefore, it is easy to add NVMe flash devices and take advantage of their high performance and low latency capabilities.



REMEMBER

While co-locating the processor and NVMe storage devices within a computer system improves performance for an application environment, the storage is captive and the data can't be readily shared across multiple computer systems. Additionally, because the bottlenecks in the storage device (by using flash) and the interconnect (by using the PCIe bus) have been reduced tremendously, now the computer's processors have become the bottleneck.

## Sharing NVMe Storage on a Network

Another way to deploy NVMe-based storage is by a shared network. Considering the advantages of co-locating NVMe storage close to the processor, one might reasonably ask, “Why would I want to place a fast NVMe storage solution behind a traditional network?” Good question.

NVMe over Fabrics (NVMeOF) is an architecture that supports a range of storage networking fabrics. These fabrics enable scaling out to large numbers of NVMe storage devices and extend the distance within a data center over which devices and subsystems can be accessed. Today, among the fabrics supported are Ethernet, Fibre Channel, and InfiniBand.

## Scaling beyond a rack

NVMeOF expands the reach of NVMe storage from something that plugs directly into one of the server's PCIe slots, to storage devices located anywhere in the data center. You can use your network as the backbone and scale your compute and storage resources across racks, across buildings, or even across metro locations. Because NVMeOF enables flexible deployment of compute and storage resources, NVMeOF storage enables many applications and application servers to share data and increase efficiencies in the data center.

## Scaling



TIP

Another key advantage of the NVMe architecture is that it supports the scaling out to large numbers of NVMe devices — theoretically in the hundreds, or even thousands.

## Data services are important

A shared storage system should provide resiliency, availability, flexibility, security, and efficiency. As an example, an enterprise-class data storage solution should ensure that data is available to the applications all the time, regardless of hardware failures in the system. It will also prevent data breaches and efficiently manage its resources.

These capabilities within a shared storage solution eliminate the need for every single business-critical application that uses the data set to provide them, in addition to the core function of the application environment. Currently, many NVMeOF storage products don't provide a full complement of these capabilities or data services. In shared storage environments, data services such as high availability, local and remote data protection, and data reduction are table-stake features.

## NVMe in Shared Storage

A third method is to deploy a shared storage solution that is connected via standard block or file protocols and uses NVMe-based flash as a storage medium. Today, many enterprise applications use traditional storage protocols such as Fibre Channel, Internet SCSI (iSCSI), Network File System (NFS), or Server Message Block (SMB) to

transfer data between compute servers and a shared storage system. Changing these application environments to use a new storage access protocol such as NVMe may not be as easy as it seems. That doesn't mean that NVMe isn't ideal for a shared storage system. Using the NVMe protocol to access solid-state storage media within a shared storage system and continuing to use traditional storage protocols to access the shared storage system provides the best of both worlds. It breaks down performance limitations within the storage system and doesn't require changes in the application environments.

With this architecture, the storage system provides all the data services required by the application environment and backend scalability of storage resources across racks, data centers, and metropolitan areas.



TIP

The IntelliFlash N-Series All-NVMe arrays from Western Digital deliver a full suite of data services and boost enterprise application performance by using Western Digital's Ultrastar NVMe SSDs for persistent storage.

## IN THIS CHAPTER

- » Looking at applications that can benefit from NVMe
- » Learning how applications can enhance a user's experience with NVMe
- » Benefitting from the speed of NVMe

# Chapter 4

## Making Your Data Come Alive

**N**on-Volatile Memory Express (NVMe) excels in a wide range of application environments. In many cases, it can boost performance and enhance user experiences. Applications that are highly transactional, or where large quantities of data must be transferred rapidly between the compute and storage environments, are good candidates. This chapter showcases a few use cases.

# Large Web-Based Platforms

People are accustomed to near-instant responsiveness from social media applications when asking for an item that must be retrieved from a massive database. Hard disk drives (HDDs) can't deliver an acceptable level of service, whereas flash storage can. As the workloads increase, connections via NVMe enable the full performance of flash to be realized, minimizing any end-user delay.

## Relational Databases

Some databases are truly massive — so much so that keeping the entire data in solid-state storage would be cost prohibitive. However, in this scenario, there's a place for NVMe-based storage to deliver write-through caches, where hot files, indexes, and metadata can be maintained for quick access. This can speed response times by an order of magnitude compared with an all-HDD storage system.

## OLTP

Classic examples of online transaction processing (OLTP) applications are airline reservation and payment processing systems. These applications require that queries (such as possible flights) are completed within seconds and can process payments with near-instant verification. Data



processing load is enormous and growing; therefore, performance must continually increase just to keep pace. Storage responsiveness is key to overall transaction speed.

## **Financial Services**

In high-frequency trading, milliseconds can mean millions of dollars. If a trader can get an order in before the competition, he can make huge profits. Paramount to success is the speed with which the decision algorithm receives relevant data. The microseconds of read latency between the storage and processor is a key part of the total.

## **Scientific Data Processing**

Scientific instruments generate huge quantities of data that must be distributed and analyzed rapidly. The Large Hadron Collider (LHC) in Europe, the site of the Nobel Prize-winning discovery of the Higgs Boson, is a prime example. Data from the LHC is distributed to 170 different sites all around the world. Low latency is essential for remote sites to leverage this data in their workflows.

## **Seismic Data Processing**

Oil and gas exploration generates large quantities of data that are sent to a storage system, and subsequently

analyzed. Even more critical is the rapid dissemination of earthquake warnings, where a few seconds can mean the difference between life and death.

## Content Caching

On the web, page load time is critical. If a page requester is physically remote from the desired content, the wait time can be hundreds of milliseconds. By caching frequently requested content on distributed servers, the closest cache will fulfill the request more rapidly than the original host.



WARNING

Today, online users expect any page they request to load in under two seconds, or they're likely to abandon the request. That abandonment can be costly, so anything that speeds the process is probably worth doing.

Caching, of course, makes just as much sense within a single data center as it does across a vast network. If content is needed as fast as possible and is requested repeatedly, it should be cached behind an NVMe interface.

## IN THIS CHAPTER

- » Determining if NVMe is right for your environment
- » Looking at NVMe storage solutions available to you

# Chapter 5

## Increasing Your Business Value by Adopting NVMe

**S**hould you take the leap into Non-Volatile Memory Express (NVMe)-based storage? This is a question that many IT directors and corporate executives are asking. A lot depends on your plan to use it and whether the adoption of NVMe enhances your potential to deliver higher revenues and more profits. To decide whether the

time is right to move to an NVMe-based storage solution, consider the following:

- » Do you have a business-critical database that's straining under heavy workloads, such as online transaction processing (OLTP)? NVMe storage can smooth out performance peaks and help break the IO bottleneck that often arises with such applications.
- » Do you have IO-intensive applications that are required to deliver real-time user experiences? If so, NVMe can help you realize the full performance that flash is intrinsically capable of providing.
- » Do you have High Performance Computing (HPC) workloads? NVMe's performance advantage can reduce total cost of ownership by a factor of four in HPC environments.
- » Is it difficult to find time to replicate or back up a database in your dynamic, real-time environment? If you can complete database tasks more quickly due to the higher performance of NVMe, you can increase data protection scheduling flexibility.
- » Do you operate virtualized server or desktop environments? NVMe-based storage can improve the performance of your multi-tenant applications, large databases, and heterogeneous workloads,

thus helping you achieve higher resource utilization and cost efficiency overall.

## **Vertical Innovation and Integration**

There are tremendous amounts of intellectual property, efficiency, and economies of scale that accrue when an organization controls every product manufacturing stage.

Western Digital, with the acquisitions of HGST, SanDisk, and Tegile, can perform every step in the production of a storage system, from flash chip fabrication to Solid State Drive (SSD) production and packaging, to integration of a storage system, complete with data management and data services. With this capability, Western Digital provides a comprehensive set of NVMe storage solutions that can increase your organization's revenue streams and profits.

### **NVMe SSDs from Western Digital**

Western Digital offers the Ultrastar SN200 series enterprise-class SSDs for infrastructure administrators who want to incorporate NVMe SSDs in their application servers. The Ultrastar SN200 stores up to 7.68 terabytes (TB), in two different form factors with enterprise grade reliability.



TIP

The SN200 series is appropriate for supporting mission critical applications that require high-performance storage close to the processors, eliminating latencies that may be induced by a fabric connected, shared storage solution. The SN200 series is ideal for data analytics and solution modeling application environments that can disaggregate workloads in HPC farms and run concurrent workloads.

## Western Digital's IntelliFlash All-NVMe storage arrays

Western Digital integrates Ultrastar NVMe SSDs into its flagship IntelliFlash storage arrays to provide an All-NVMe unified, flash storage system. The IntelliFlash N Series All-NVMe arrays deliver high performance at very low latency along with nearly 700TB of effective capacity in a compact 2 rack unit (RU) footprint.

IntelliFlash N Series All-NVMe arrays provide a complete set of data management and data services including inline deduplication and compression, snapshots, read/write clones, and thin provisioning, as with the T Series All-Flash arrays. Cost-effective disaster recovery solution can be achieved by replicating data between IntelliFlash All-NVMe arrays and All-Flash or Hybrid Flash arrays.

- » Blazing a trail for NGPM
- » Embracing a hybrid cloud solution

## Chapter 6

# The Future of NVMe Storage

The migration to Solid State Drives (SSDs) using flash as primary storage devices represents a major break from the dominance of rotating media. The Non-Volatile Memory Express (NVMe) architecture that's designed to support solid-state devices, as well as future technologies that are currently in development, enables the movement of large amounts of data at very low latencies. What if you were able to persist large amounts of data at an order of magnitude lower latency?

# NGPM Technologies

Flash storage is blazing a trail for next-generation persistent memory (NGPM) such as Phase Change Memory (PCM) and 3D-XPoint, currently under development. These new persistent memory technologies, also known as *Storage Class Memory* (SCM), will feature latencies that are far lower than present day NAND flash devices. They can make use of the NVMe interface to minimize the total end-to-end latency of the processor to storage path. Imagine running entire applications in-memory while providing data persistence.

## Persistent Memory on a Fabric

Because NVMe works over fabrics (networks) such as Fibre Channel and Ethernet, the deployment scale of a storage environment goes beyond a rack to across data centers and metro locations. This fabric-based scale may allow many organizations to fully embrace a hybrid cloud solution (a combination of public and private cloud) and provide the flexibility and agility of deployment with service levels for performance, availability, and security.



## IN THIS CHAPTER

- » Realizing the full benefits of NVMe storage
- » Making your applications faster

# Chapter 7

## Ten NVMe Storage Takeaways

In this chapter, you get ten takeaways about Non-Volatile Memory Express (NVMe)-based storage:

- » New, high-performance persistent storage technologies require a new, high-performance interface for the full benefits to be realized.
- » The NVMe interface is a greenfield design, not hampered by backward compatibility for legacy devices.

- » The NVMe interface is designed to interconnect with Peripheral Component Interconnect Express (PCIe)-compatible solid-state storage devices.
- » NVMe supports parallel operations of up to 65,535 queues to eliminate transmission bottlenecks.
- » NVMe support 65,536 commands in every queue, which means that you could in theory line up more than 4 billion commands for action.
- » NVMe can deliver twice as many Inputs-Outputs Per Second (IOPS) per core as Serial Attached SCSI (SAS) or Serial ATA (SATA) interfaces.
- » Hot-swapping of NVMe storage is supported in the 2.5-inch form factor devices.
- » Non-Volatile Memory Express Over Fabrics (NVMeOF) extends the benefits of NVMe storage out into the entire data center.
- » A shared storage system using traditional storage protocols can utilize NVMe-based Solid State Drives (SSDs) to deliver very high performance at low latencies, without overhauling the entire application environment.
- » A number of real-time workloads that require high performance at very low latency can benefit from the fast IO transfers between processors and storage made possible by the NVMe interface.

# Make environments for applications to thrive

Creating environments for data to thrive requires a deep understanding of how it's transforming everything from research to banking to personal mobility. The next few decades will usher in a future that's more predictive, more productive, and more personal. Adopt an intelligent storage solution that understands exactly what it needs to thrive — wherever it's needed.

## Inside...

- Harness NVMe to make your applications thrive
- Enable your business to transform and deliver more profits
- Extract information from your data

Go to [Dummies.com](http://Dummies.com)<sup>®</sup>  
for videos, step-by-step photos,  
how-to articles, or to shop!

**Western Digital.**

**for  
dummies**<sup>®</sup>  
A Wiley Brand



Also available  
as an e-book

ISBN: 978-1-119-50971-4

Not for resale

# **WILEY END USER LICENSE AGREEMENT**

Go to [www.wiley.com/go/eula](http://www.wiley.com/go/eula) to access Wiley's ebook EULA.