

Making Everything Easier!™

Tintri Special Edition

VM-Aware Storage

FOR
DUMMIES®
A Wiley Brand

Brought to you by



Tintri

Lawrence C. Miller, CISSP



About Tintri

Tintri VM-aware storage is the simplest for virtualized applications and cloud. Organizations including Toyota, United Healthcare, NASA, and seven of the Fortune 15 have said “No to LUNs.” With Tintri they manage only virtual machines, in a fraction of the footprint and at far lower cost than conventional storage. Tintri offers them the choice of all-flash or hybrid-flash platform, converged or stand-alone structure, and any hypervisor. Rather than obsess with storage, leaders focus on the business applications that drive value — and that requires that they keep storage simple.

For more information, visit www.tintri.com and follow Tintri on Twitter: @tintri.

VM-Aware Storage

FOR
DUMMIES[®]
A Wiley Brand

Tintri Special Edition

by Lawrence C. Miller, CISSP

FOR
DUMMIES[®]
A Wiley Brand

VM-Aware Storage For Dummies®, Tintri Special Edition

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

Copyright © 2016 by John Wiley & Sons, Inc., Hoboken, New Jersey

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

Trademarks: Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. 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 978-1-119-28856-5 (pbk); ISBN 978-1-119-28857-2 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgments

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

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

Project Editor: Martin V. Minner

Acquisitions Editor: Amy Fandrei

Editorial Manager: Rev Mengle

Business Development Representative:
Karen Hattan

Production Editor: Siddique Shaik

Tintri Contributors: Lei Yang,
Craig Schultz, Sunil Prasad
Muralidhar, Vijay Iyer, Tomer Hagay,
Christopher Slater, Neil Glick,
Vineet Kakani, Yael Zheng,
Scott Buchanan, Brian Houlihan

Table of Contents

.....

Introduction	1
About This Book	1
Foolish Assumptions	2
Icons Used in This Book.....	2
Beyond the Book.....	2
Chapter 1: Seeking Innovative Alternatives to Conventional Storage	3
Virtualization Is Real and the Cloud Isn't on the Horizon — It's Here!.....	3
Conventional Storage Is a (Real) Virtual Drag.....	5
Prepare for the Cloud Now!	9
Chapter 2: Discovering How VM-Aware Storage Can Help You	11
Understanding Changing Storage Requirements	11
VM-Aware Storage Basics	13
Solving Modern Storage Challenges with VM-Aware Storage.....	17
Scaling out to sustain growth.....	18
Gaining insight into poor storage performance.....	18
Eliminating waste and guaranteeing performance....	19
Reducing management complexity	19
Chapter 3: Maxing Out the Value of VM-Aware Storage	21
Flash! A-ah It Isn't (Quite) the Savior of the Storage Universe	21
Got Virtual Apps? Guarantee Their Performance with QoS	23
I Can See Clearly Now . . . with Predictive Analytics	25
Time Is Money — Scaling Out While Saving Both.....	27

Chapter 4: Exploring VM-Aware Storage Use Cases**31**

- Hitting Critical Mass with Business-Critical Applications31
- Enjoying the Benefits of the Cloud (On-premises and Hosted) 33
- Empowering Your Workforce with VDI (That Actually Works) 36
- Benefiting from DevOps with Agile and Grace 38

Chapter 5: Ten Cool Things You Can Do Only with VM-Aware Storage**41**

Introduction

There's a disconnect in your data center. You've virtualized your applications, and your teams are thinking and working in virtual machines (VMs). But your storage hasn't materially changed in decades — it still uses the same constructs (such as *logical unit numbers*, or LUNs) that were introduced for physical server workloads during an era that has long since passed. That disconnect between virtualized applications and conventional storage is costing your business time and money.

Many storage innovations are happening, and you might be exploring or deploying some of them to address the storage pain in your data center. But it doesn't matter if those products aren't VM-aware. A big claim? Yes, but it's an important concept to understand.

Only VM-aware storage (VAS) is specifically built for virtualized applications, stripping out the complexity of physical storage management so that you can focus on the applications that matter most for your business. With VAS, you can manage storage functions on individual applications, saving you time, money — and sanity.

About This Book

Although many storage vendors claim the “VM-aware” label, few are truly VM-aware. Here's a hint: If the product's architecture is built on a foundation of storage LUNs and volumes, it can never be fully VM-aware. If you want to learn more, keep reading! *VM-Aware Storage For Dummies* will help you understand what VM-aware storage truly is and what it can do for your virtualized and cloud environments.

Foolish Assumptions

It's been said that most assumptions have outlived their usefulness, but I'll assume a few things nonetheless! Mainly, I assume that:

- ✔ You're a manager or team member in an organization that is planning a virtualization and data storage project.
- ✔ You aren't necessarily a technical expert, but you'll benefit from an understanding of VM-aware storage concepts and technologies while navigating the overcrowded and confounding array of storage solutions that are available to you.

Icons Used in This Book

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



You won't find a map of the human genome here, but if you seek to attain the seventh level of NERD-vana, perk up! This icon explains the jargon beneath the jargon!



Tips are always appreciated, never expected — and you'll certainly appreciate the helpful suggestions, soul-searching advice, and handy nuggets of information this icon points out.



You should commit this important information to your non-volatile memory or your gray matter — it could save you some time and frustration!



This icon points out the stuff your mother warned you about. Okay, probably not. But you should take heed nonetheless to avoid costly mistakes or pitfalls — and alligators.

Beyond the Book

There's only so much I can cover in 48 short pages, so if you find yourself at the end of this book, thinking "Gosh, this was an amazing book; where can I learn more?" just go to www.tintri.com.

Chapter 1

Seeking Innovative Alternatives to Conventional Storage

.....

In This Chapter

- ▶ Appreciating the real benefits of virtualization and the cloud
 - ▶ Understanding conventional storage challenges
 - ▶ Embracing (hugging) a cloud strategy that is right for you
-

This chapter gets you up to speed on virtualization technology, and it helps you understand the challenges and limitations of conventional storage for the virtual and cloud environments.

Virtualization Is Real and the Cloud Isn't on the Horizon — It's Here!

In case you've spent more time keeping up with the Kardashians than keeping up on technology, let me quickly bring you up to date: Virtualization has been *the* hottest technology over the past decade!

Storage 101

Before you dive headlong into VM-aware storage, let's cover a few basic storage terms and concepts:

- ✔ **File-based storage** organizes data into files and directories within a hierarchical directory-based structure.
- ✔ **Block-based storage** organizes data as a sequence of bits or bytes of a fixed length or size. To use block storage, you must first create a volume or **logical unit number (LUN)** — a logical abstraction or virtualization layer between a physical storage device/volume and applications.
- ✔ A **hard disk drive (HDD)** stores data using different magnetic patterns on a spinning disk. HDD performance is typically hundreds of input/output operations per second (IOPS) with a latency of milliseconds.
- ✔ A **solid-state device (SSD)** stores data electrically in non-volatile memory using NAND flash technology. SSD storage performance is typically tens of thousands to millions of IOPS with a latency of microseconds.
- ✔ **Data management** refers to various storage operations, including clones, replication, and snapshots. A **clone** creates a copy of a storage volume. **Replication** duplicates data in real time to another physical storage device. A **snapshot** consists of an initial point-in-time clone of a storage volume and subsequent copies of only the changes that have occurred in the storage volume since the initial clone.



A *hypervisor* is *virtualization* software that runs between the hardware and the operating system(s) running *virtual machine* (VM) workload(s). The hypervisor enables multiple independent VMs or applications to run on a single physical server, by abstracting the resources (processors, memory, networking, and storage) of the physical server.

Virtualization has transformed the modern data center and is a key enabling technology for the cloud environments. In the not-too-distant past, individual enterprise applications and databases were installed on dedicated physical servers. But deploying an application on its own physical server hardware has some important technical limitations beyond cost effectiveness that negatively affect business agility and efficiency, including:

- ✔ **Poor scalability:** As application usage and functionality increase because of growing business demands and feature upgrades, application performance tends to degrade. The typical answer to poor performance and capacity limitations — adding more powerful servers, more storage, and faster networking — is costly, time-consuming, and fraught with risk.
- ✔ **Underutilization of resources:** Even the most resource-intensive applications rarely take full advantage of today's powerful server hardware, storage, and networking. Inefficient usage results in wasted capacity and is also costly in terms of data center space, power consumption, and cooling.



Virtualization enables server consolidation, which greatly improves application scalability and data center resource utilization. VMs can be quickly provisioned for new or upgraded applications and easily migrated to different physical servers and geographic locations. A single physical server can run multiple VMs and virtualized applications to better utilize resources.

Cloud computing extends virtualization benefits and enables businesses to deploy virtualized applications at massive scale (see Figure 1-1).




Figure 1-1: Physical, virtual, and cloud computing architectures.

Conventional Storage Is a (Real) Virtual Drag

The increased use of server virtualization in data center environments has dramatically improved the efficiency of many IT processes, including data storage provisioning.

Throughout this book, I look at various attempts by storage vendors to place conventional filesystem or block storage in various parts of the data center architecture in an effort to share the storage pool more effectively for virtualized applications. However, conventional approaches weren't designed for virtualization and have proven to be a forced fit at best for today's highly virtualized data centers.

While DAS- and SAN-based storage platforms present data as blocks, and NAS-based storage manages data as files (which are closest to the nature of VMs, or VMDK files, but still not VM-aware), all three have no concept of virtual machines, creating various layers of hardware and management complexity.

What these conventional storage architectures lack is a more granular and flexible unit of management for storage. In order to effectively manage virtualized workloads, you need a VM-aware storage solution or filesystem.

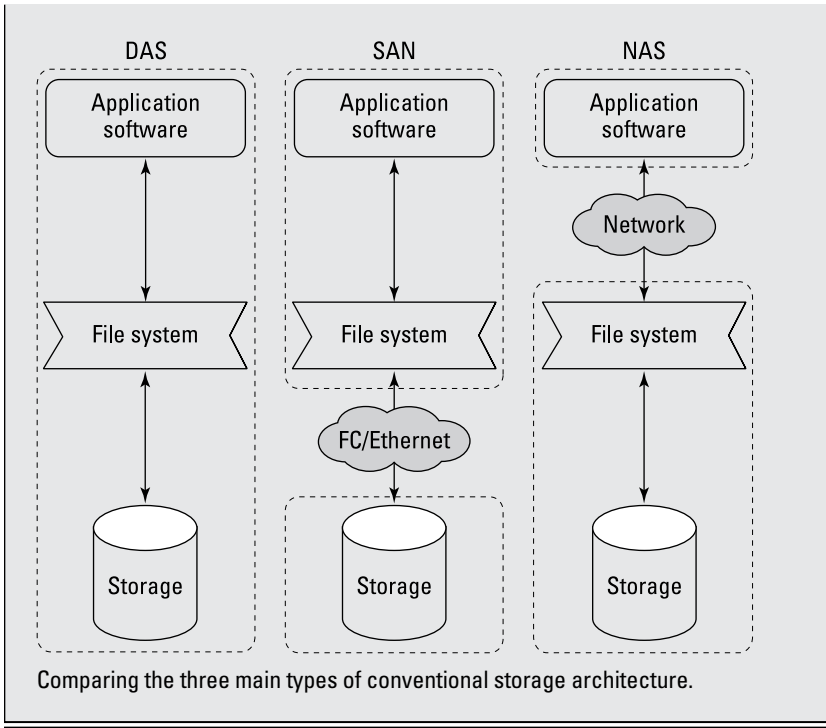
Conventional wisdom on storage architecture

Conventional storage architectures (see the figure) include:

- ✔ **Direct-attached storage (DAS)** is data storage either physically located inside or directly attached to a server, typically using SATA, Fibre Channel, or some SCSI variant, and commonly used in smaller application servers or for working with very small data sets that do not require directly sharing with other servers.
- ✔ A **storage area network (SAN)** is a dedicated — and expensive — storage network. A SAN consists

of storage arrays, controllers, storage routers/switches, storage management software, and specialized network cables and connectors.

- ✔ **Network-attached storage (NAS)** is a data storage appliance that is typically connected to servers through a network connection, such as Ethernet, and unlike DAS or SAN, provides a file-based, rather than block-based, interface to an operating system. NAS typically communicates with the server using the network file system (NFS) protocol.



In a physical storage architecture, it's straightforward and simple to map storage in a 1:1 relationship with physical servers hosting individual applications (see Figure 1-2).




Figure 1-2: Physical applications mapped to physical storage in a 1:1 relationship.

Then virtualization comes along, which enables multiple VMs and applications to share a common storage pool. Conventional storage platforms were built and deployed before virtualization became a defining characteristic in the modern data centers, and thus weren't designed to support virtualized workloads. Conventional storage architectures have limitations that restrict their effectiveness in virtualized environments, including:

- Poor application performance visibility:** Before server virtualization, conventional block storage primitives could be mapped to a monolithic application installed on a single physical server in a 1:1 relationship. However, in a virtualized environment, an I/O “blender effect” (see Figure 1-3) is created — operations from multiple applications running as VMs on a single physical server are mapped to a conventional storage container, such as a LUN or volume.

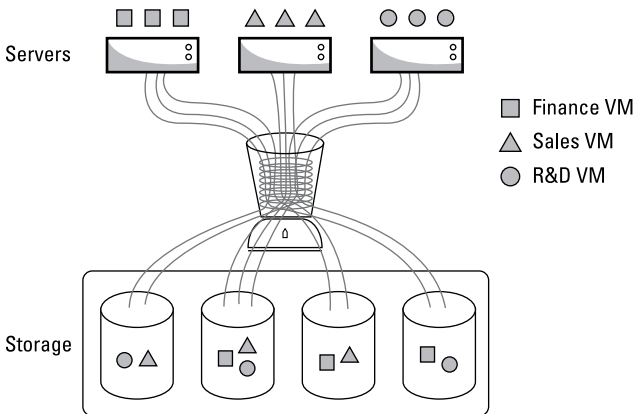


Figure 1-3: Mapping virtualized application VMs to conventional LUN-based storage is complex.



The lack of visibility created by the I/O “blender effect” makes understanding and troubleshooting storage capacity and performance issues extremely challenging (like trying to put the ingredients of a blended fruit smoothie back together to figure out if it was the strawberries, bananas, mangos, or jalapeños that didn’t quite taste right).

- ✔ **Waste resulting from over-provisioning:** Virtualized environments tend to evolve quickly. To meet performance requirements and application infrastructure needs, IT system designers often buy more storage capacity and storage tiers than they need — a process called *over-provisioning*. Over-provisioning is incredibly inefficient in terms of cost and space; much of the storage capacity and performance may not be needed for months or even years, but it has to be paid for up front.
- ✔ **Management complexity:** Organizations create new VMs, move workloads, and assign virtualized compute resources using management tools designed for virtualized environments. However, conventional storage — built on block storage primitives as the unit of management — is not well suited to management processes that have become commonplace within virtualized environments.



For example, in some instances, virtualization enables workloads to scale up or down automatically, based on demand, but most conventional storage technologies require manual intervention to assign the corresponding storage capacity before the workload can be started.

Prepare for the Cloud Now!

In *As You Like It*, Shakespeare extols the virtues of moderation, asking “Can one desire too much of a good thing?” But we’re talking about server virtualization — not chocolate or exercise — and Shakespeare didn’t know squat about server virtualization! So if your organization has embraced virtualization, how do you get ready for virtualization at scale? By deploying a cloud-ready infrastructure.

And this is when things get complicated with conventional storage and its associated manual and error-prone storage management processes. These storage management challenges, which you may have painstakingly lived with in smaller virtualized environments, become insurmountable with virtualization at scale. Again, Shakespeare didn’t understand the management complexity of virtualization at scale when he wrote “All the world’s a stage . . . and one man in his time plays many parts.” Many parts yes, but several thousand

parts (or manual storage management processes)? Get real! These “parts” include:

- ✔ **Monitoring and management:** Data protection and quality of service (QoS) policies for thousands of individual VMs and virtualized applications must be visible and persistent across different, geographically dispersed, hosted (multiple cloud service providers), and on-premises (multiple data centers) cloud infrastructures.
- ✔ **Automation:** Storage infrastructure must be customizable to support automation of storage management tasks at scale, but at a granular, individual VM, or application (through a service group) level. This means enabling broad application programming interface (API) support for developers, cloud service providers, large enterprises, and other automated virtualization technologies.
- ✔ **Agility:** The cloud enables business agility and faster time to market with self-service IT providing the ability to spin up and tear down thousands of VMs on demand — but only if it’s safe, intuitive, and quick to deploy! A series of manual and error-prone steps will only slow the business down.

Chapter 2

Discovering How VM-Aware Storage Can Help You

.....

In This Chapter

- ▶ Understanding the need for VM-aware storage
 - ▶ Learning the basics of VM-aware storage
 - ▶ Tackling modern storage challenges with VM-aware storage
-

In this chapter, you learn how storage has changed in today's virtualized world, what VM-aware storage is, how it works, how it solves modern storage challenges, and how it addresses conventional storage limitations for the virtual and cloud environments.

Understanding Changing Storage Requirements

Before applications were virtualized, tuning storage to support different applications was a complex, but well understood, process. Generally speaking, each application ran on a dedicated server and on storage hardware with capacity and performance characteristics matched to each application's unique requirements. However, this approach was inefficient.

With virtualization, multiple applications running as VMs share storage, making it (virtually) impossible to specifically

allocate dedicated storage hardware to meet the required capacity and performance characteristics of each application. However, the need for “dedicated” storage for specific applications still exists — you still need to match the right storage performance characteristics to each individual application. With conventional storage, this requires constant tuning of logical unit numbers (LUNs) and volumes that contain numerous VM workloads — all with different storage performance requirements.

Table 2-1 gives a few examples of the varying storage performance requirements that might exist in different application use cases.

Table 2-1 Examples of storage performance requirements for different application use cases

	<i>Read performance</i>	<i>Write performance</i>	<i>Random access performance</i>	<i>Latency</i>	<i>Concurrent users</i>
Backing up critical data files	Low	High	Not applicable	Not applicable	Low
Finding a record in a database	High	Low	High	Low	Moderate
Starting up many virtual desktops	High	Low	N/A	Moderate	High



Other potential storage characteristics to consider for different applications include:

- ✓ Amount of data redundancy
- ✓ Block size
- ✓ Cost per unit of capacity
- ✓ Data protection mechanisms and protection schedules

- ✔ File transfer protocol support
- ✔ Potential for data reduction
- ✔ Reliability

A valid question you should be asking now is: How can you possibly match the different characteristics and requirements in conventional storage?

VM-Aware Storage Basics

VM-aware storage restores simplicity to the storage system in virtualized and cloud environments by getting rid of LUNs and other conventional storage constructs on the back end. VM-aware storage also abstracts the complexity of redundant array of independent disk (RAID) configurations under the hood from the daily tasks of storage administrators. With VM-aware storage, you can focus on applications or VMs and re-create a virtual 1:1 mapping between your applications and storage — the I/O “un-blender” (see Figure 2-1)!




Figure 2-1: VM-aware storage maps storage performance characteristics to virtualized application VMs in a 1:1 relationship.

VM-aware storage has built-in intelligence about the virtualized applications that use it. This intelligence can include a number of factors, such as an understanding of:

- ✔ Acceptable delays in finding data (latency)
- ✔ Data availability

- ✓ Data read/write patterns
- ✓ How to protect applications and data in various failure scenarios

VM-aware storage works with the virtualization infrastructure to understand the performance needs of individual virtualized applications and make the storage system more responsive to those needs.

VM-aware storage eliminates complexity in storage management tasks by acting as a dedicated resource manager, delivering the right level of performance to each virtualized application.



VM-aware storage is like being able to observe all vehicles traveling through a particular town every day and knowing where they started from and where they are going, as well as the individual driving habits and abilities of each driver. Using this information, you can organize the traffic light patterns across town to make each vehicle's journey flow smoothly, and you can change the direction of reversible lanes during rush hour — virtually eliminating traffic jams forever!

VM-aware storage optimizes how data is stored, based on its understanding of virtualized applications and their requirements. It uses storage technologies, such as deduplication and compression, to improve storage efficiency, and quality of service (QoS) policies to meet different application performance requirements. These capabilities enable you to independently manage and optimize storage for individual applications.



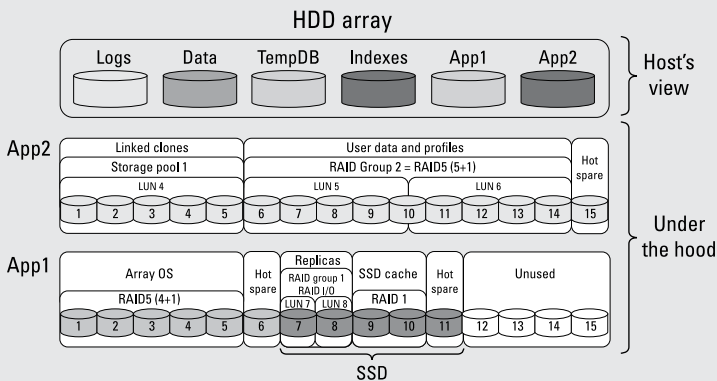
VM-aware storage creates *virtual disks* (vDisks), which a VM uses for its operating system, application software, and other data files. A vDisk hides the physical storage resources from the VM's operating system in much the same way that a hypervisor hides the physical server resources from the VM's operating system. Regardless of the type of physical storage device, the vDisk always appears to the VM as local storage.

At this point you may be thinking, “Wait, VM-aware storage does all of this without LUNs?” Yes! Read the sidebar “Look Mom, no LUNs!” to learn how.

Look Mom, no LUNs!

Remember the thrill you used to get from causing panic and terror in your mother when you rode your bicycle without holding the handlebars? Now, you can relive those carefree days and create that same feeling of panic and terror in your storage administrators — without the recklessness and danger — when you deploy VM-aware storage and yell, “Look Mom, no LUNs” (although it may be a bit awkward calling your storage administrator “Mom”)!

How can you get away without LUNs? In conventional storage systems, storage performance characteristics and storage operations are assigned and executed at the LUN or volume level (see the figure). Because numerous VMs and applications with different performance and data protection requirements are placed on a single LUN or volume, conventional storage lacks the ability to provide granular control of important features and functions.



Conventional storage uses LUNs and volumes to define a filesystem, and applies RAID configurations at that level.

The lack of granularity in the filesystem of conventional storage systems creates various storage management issues, including:

- Inflexible allocation of performance and capacity


- Wasteful over-provisioning to ensure peak and future performance needs
- Manual processes to reclaim space and continually tune the storage array as workloads change

(continued)

(continued)

VM-aware storage defines the file-system at the VM-level and provides a thin/thick provisioning layer to abstract the complexity of how and where data is stored (see the

figure). The unit of management is simply the VM, and storage is treated as a datastore with thin/thick provisioning.



VM-aware storage removes the complexity of LUNs and RAID configuration and defines the filesystem characteristics at the individual VM level.



VM-aware storage enables numerous storage management advantages over conventional storage, including:

- Capacity is automatically recaptured and made available to other VMs as workloads are scaled back.
- Protection schemas can be based on data type or even individual VMs, and these policy-based definitions can adapt to new storage in a non-disruptive fashion.
- Thin provisioning means that adding capacity doesn't require reconfiguration of VMs or applications, hence no downtime.

If you aren't convinced yet about the power of doing away with LUNs, think about replication as an example.

In conventional storage platforms, storage operations such as replication and snapshotting are inefficient because they must be performed at the volume or LUN level. Figure 2-2 illustrates an example of a production VM (VM1) that an organization needs to replicate to a remote site. On a conventional storage system (the top part of the figure), other VMs located on the same LUN (VMs 2 through 5) are also replicated to a

remote site over the wide area network (WAN), wasting processing power to perform the cloning and replication, time, bandwidth, and storage capacity. In contrast, on a VM-aware storage system (the bottom part of the figure), only VM1 is replicated across the WAN, hence no hitchhikers.

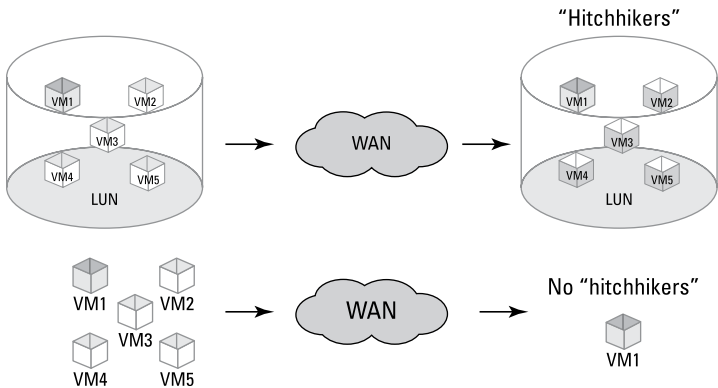


Figure 2-2: Replication over WAN in conventional storage performed at the LUN level versus VM-aware storage at the VM level.

Because the filesystem in VM-aware storage is defined at the individual VM level, storage operations — such as snapshotting and replication — can be performed at the individual VM or application level. This results in more granular and efficient operations.

Solving Modern Storage Challenges with VM-Aware Storage

VM-aware storage solves modern storage challenges in virtualized and cloud environments (discussed in Chapter 1), including:

- ✔ Limited scalability
- ✔ Lack of visibility
- ✔ Over-provisioning (and lack of QoS)
- ✔ Management complexity

Scaling out to sustain growth

Accurately predicting how much data storage capacity your organization will need in the next three months is incredibly challenging, let alone over the next three to five years — you'd probably have better luck forecasting the weather over the next three months! Further complicating this challenge is the fact that storage technology changes as often as, well — the weather. This dilemma often results in organizations adding too much complexity and creating too many islands of storage that don't integrate with other solutions in your data center, hence having limited scalability.

VM-aware storage leverages a *scale-out* architecture that enables you to optimize the placement of workloads over time, scaling both storage capacity and performance independently. New storage appliances are added into a datastore pool so that the collective VMs and applications can access the right storage without requiring reconfiguration.

Gaining insight into poor storage performance

The I/O blender effect (discussed in Chapter 1) creates a blind spot in virtualized applications running on conventional storage. The lack of visibility makes troubleshooting a time-consuming, costly, and frustrating trial-and-error process that leaves storage administrators guessing about the root cause of storage performance issues, while business users are left unproductive and idling.

VM-aware storage uses a model, built around the requirements of the applications that use it, to shape an appropriate set of storage resources (such as flash, compression, deduplication, and backup and recovery procedures). Hence, it can help to build the usage model and validate that the storage system is meeting the needs of the virtualized applications.

VM-aware storage gathers and presents information at the individual VM level. It tracks all the characteristics of resident VMs and can correlate with the rest of the virtualized infrastructure (host, network, and storage) in real-time. Thus, administrators are better able to spot issues and resolve

performance problems in much less time than with conventional storage systems.

Eliminating waste and guaranteeing performance

In conventional storage systems, over-provisioning is a common practice to ensure that applications always have sufficient storage capacity and performance. This well-intentioned but misguided practice causes under-utilization of storage capacity and limited ability to guarantee performance.

VM-aware storage intelligently utilizes several technologies at a granular VM level, as opposed to using a “shotgun” approach on a storage LUN, including:

- ✓ **QoS** is used to guarantee performance for applications at different service levels, which are typically measured in latency and input/output operations per second (IOPS).
- ✓ **Deduplication** identifies redundant repetitive byte patterns and replaces them with a much smaller pointer to the original copy of the duplicate data pattern.
- ✓ **Compression** uses mathematical techniques to reduce the overall size of data stored on a storage system.

Reducing management complexity

In conventional storage, manual processes define how you manage data (the *control plane*) and where that data physically resides (the *data plane*). Adding or removing resources requires many repetitive “re-wire” steps, and you have to tell the operating system and installed application(s) the precise characteristics of the underlying storage.

In *software-defined storage* architectures, the control plane is built into software to abstract many complex manual storage processes into software-based solutions.

VM-aware storage uses many of the same principles as software-defined storage, but takes them a step further. For example, the storage system’s control plane can use information gleaned from the VMs. Based on a VM reporting certain threshold values (such as performance, available capacity,

and latency), it might trigger automated provisioning or make more resources available.



Many storage vendors today claim that their storage products are “VM-aware.” But what, exactly, makes VM-aware storage “VM aware”? The five requirements for VM-aware storage to truly be VM-aware are:

- VM-level filesystem
- VM-level storage operations
- VM-level automation
- VM-level QoS
- VM-level analytics

Is your storage VM-aware? Take this simple test to find out

Many storage vendors are jumping on the VM-aware storage (VAS) bandwagon and now tout their products as “VM-aware,” “VM-centric,” or some other marketing catchphrase. But it takes more than a creative label for a storage platform to truly be VM-aware. It’s fairly simple to tell when a storage product isn’t truly VM-aware. Simply answer the following question:

Are there any LUNs or volumes in the storage architecture?

If the answer is YES, do not pass “Go,” do not collect \$200. You’re looking at conventional storage — it is not VM-aware, no matter what the label says!

If the answer is NO, it *may* be VM-aware storage. A further test kit is:

Do they meet the five requirements of VM-aware storage?

A VM-level filesystem is the first requirement of truly VM-aware storage, followed by VM-level storage operations, VM-level automation, VM-level QoS, and VM-level analytics.

Achieving any of these requirements is not easy because it requires a ground-up redesign instead of retrofitting on conventional architecture.

Don’t worry if you can’t readily recall the five requirements of VM-aware storage. I call them out for you as I cover them throughout this book!

Chapter 3

Maxing Out the Value of VM-Aware Storage

In This Chapter

- ▶ Recognizing the benefits and limitations of flash storage technology
- ▶ Prioritizing and guaranteeing the success of mission-critical applications
- ▶ Seeing through the fog of virtualization with predictive analytics
- ▶ Simplifying management at scale

In this chapter, you explore the business benefits of VM-aware storage features and capabilities that enable a broader virtualization and cloud strategy for your enterprise.

Flash! A-ah It Isn't (Quite) the Savior of the Storage Universe

Flash storage technology, in general, provides significantly higher performance and reliability than mechanical hard disk drive (HDD) technology (see the “Storage 101” sidebar in Chapter 1). However, flash in and of itself can't solve all of an organization's storage performance and capacity woes.

Flash technology is often used with other technologies (see Figure 3-1) within storage (both conventional and VM-aware) and IT architectures to improve performance, including:

- ✓ **Flash storage array:** In a flash storage array, solid-state devices (SSDs) are used for storage. This configuration typically provides a substantial performance increase,

but in some cases (particularly in conventional storage platforms), the effectiveness of flash technology may be limited by the underlying unit of management and handling of noisy neighbor problems.

- ✓ **Active working set and I/O profiling:** A detailed monitoring, analysis, and profiling of storage I/O patterns is performed before bits are moved to different storage media. These storage I/O patterns can be quite different and dynamic depending on the applications. VM-aware storage takes the intelligence to the VM level to help optimize these operations.
- ✓ **Inline deduplication and compression:** The resulting data reduction from these technologies is highly workload dependent. For example, virtual desktops can result in very high data reduction, while SQL databases can be highly distinct, hence low data reduction exists with the same algorithm.




Figure 3-1: Flash technology is used with other technologies, such as inline deduplication and compression, to improve storage performance.



Using flash storage technology in conventional storage that isn't VM-aware is like driving a Formula One racing car without a steering wheel. Flash storage technology runs into the same visibility, management complexity, and scalability issues as HDD storage technologies — they're just pushed out to a later time, at best. And when you get stuck in a traffic jam

(with other “noisy” VMs that happen to be on the same logical unit number, or LUN), the extra speed you *hope* to get from your Formula One car won’t get you there any faster.

Got Virtual Apps? Guarantee Their Performance with QoS

Not all applications are created equal. Some applications are business critical and must always perform consistently. For cloud service providers, service-level agreements (SLAs) often define penalties for failure to deliver a specified level of performance to your customers. For enterprises, downtime means lost business or costly errors in critical databases.

Conventional storage defines quality of service (QoS) policies at the LUN/volume level. Each LUN/volume may have dozens or hundreds of VMs with very different performance requirements, but you must assign the same QoS level to all the VMs within the LUN/volume. The result is a first-in/first-out (FIFO) funnel effect that creates I/O traffic jams, as shown on the left side in Figure 3-2. A single VM can effectively “block” all other VMs because QoS can’t be applied at a level granular enough to control “noisy” applications and guarantee appropriate performance levels.




Figure 3-2: Conventional storage systems lack granular, VM-level QoS control, which creates an I/O traffic jam.

As shown on the right side in Figure 3-2, VM-aware storage allows you to create granular QoS policies at the VM level. Each individual VM or vDisk gets its own “lane” with the

appropriate “width” (corresponding to QoS level) to isolate performance so that there are no conflicts with noisy neighbors and performance levels are guaranteed.



VM-aware storage enables organizations to define and enforce QoS policies that prioritize individual application and VM workloads, in order to guarantee minimum (or maximum) levels of performance. For example:

- ✔ **It’s Cyber Monday and your e-commerce web application is in a turkey coma.** VM-aware storage looks at what else is running and, noticing that a non-essential data archive job could be allocated fewer resources, it diverts the freed-up surplus resources to the e-commerce application to maximize sales.
- ✔ **It’s month-end and your Finance department is screaming because its finance systems aren’t.** VM-aware storage can temporarily reduce the resources available to other applications during critical month-end closing and make a believer out of your CFO.
- ✔ **It’s 8 a.m. and the tardy bell is ringing for your school’s virtualized desktops.** Boot storms are a common issue in virtual desktop infrastructure (VDI) environments (see the use case in Chapter 4). VM-aware storage can keep a copy of your desktop startup image files in flash each morning between 7:30 a.m. and 8:30 a.m., then release these files to prioritize other workloads.

VM-aware storage is designed to automate many of the manual tasks needed to define, manage, and maintain QoS at an individual VM (or groups of VMs) level, including:

- ✔ Immediate visualization of contention, performance troubleshooting, and impact of QoS setting changes
- ✔ Auto-configuration and performance optimization of workloads across multiple VM-aware storage appliances and datastores, which are later defined as a storage pool
- ✔ Historical views of VM-level storage performance that enable organizations to meet SLAs and create reports for capacity planning and trend analysis
- ✔ Integration with hypervisor management tools to greatly reduce the number of management steps required



One of the five key requirements of VM-aware storage (see the spot guide at the end of Chapter 5) is VM-level QoS.

I Can See Clearly Now... with Predictive Analytics

One of the fundamental benefits of virtualization is the ability to simplify the creation, deployment, and ongoing management of virtualized applications with software-based tools. Virtualization platform vendors such as VMware, Microsoft, Red Hat, and Citrix provide great tools to monitor, manage, and troubleshoot virtual environments. However, without the same capability on the critical data storage layer underlying these environments, you only have partial visibility, and hence, only a partial toolkit.

Many conventional storage systems provide performance averages and correlations at a LUN or volume level, but such information is not helpful for managing and troubleshooting performance in individual VMs and applications, as shown in the top diagram in Figure 3-3.

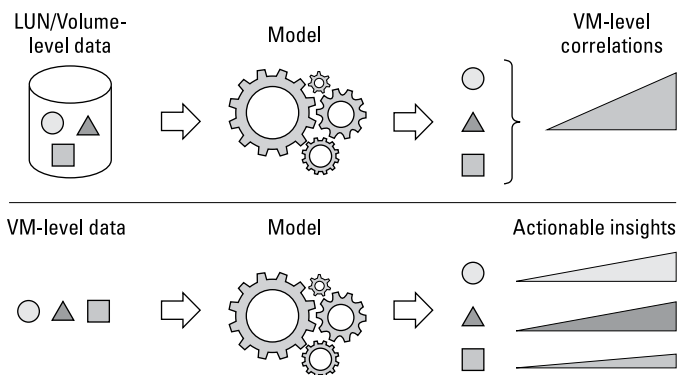


Figure 3-3: Conventional storage analytics can only provide averages and correlations, rather than the actionable insights available in VM-level analytics.

VM-level analytics take the guesswork out of troubleshooting and fine-tuning storage performance in virtualized environments. VM-aware storage provides real-time measurements

and actionable insight, including input/output operations per second (IOPS), throughput (in Mbps), and latency at a per-VM level.



The difference between conventional storage correlations and VM-level analytics can be easily understood with the following analogy: If you want to know the temperature in a room of your house, do you analyze complex modeling of other people's home temperatures and global weather trends to get an average? No, you simply read a thermometer in that room. VM-level analytics provide direct data from the VMs. Even if you can get the average, it is a few hours old at best, which doesn't meet your need for real-time information.

VM-aware storage provides the following VM-level analytics benefits over conventional storage platforms:

- ✓ VM-aware storage lets you focus on the needs of individual VMs or applications without forcing you to create “one-size-fits-all” logical unit numbers (LUNs) or volumes.
- ✓ With VMs as the universal unit of management on a VM-aware storage platform, troubleshooting issues at the host, network, and storage layers is a breeze.
- ✓ Understanding the performance characteristics of the VMs in your storage environment helps you understand your applications, which means you can talk to your virtualization and application administrators in their language.
- ✓ In-depth knowledge of your VMs and applications needs, and how they have evolved, helps you understand your future storage performance and capacity needs. Predictions are based on actual measurements — not averages or correlations — which means they will be more accurate. This eliminates over-provisioning and under-utilization in your storage infrastructure, leading to higher efficiency in your virtualized data center and cloud environment.

When future storage growth predictions are based not only on your VM's storage growth, but also on real-time performance and I/O trends, you can eliminate performance bottlenecks without piling on unnecessary resources.



One of the five key requirements of VM-aware storage (see the spot guide at the end of Chapter 5) is VM-level analytics.

Time Is Money — Scaling Out While Saving Both

Managing a few VMs across a couple of storage arrays might seem relatively simple. But as your business grows, so does your storage footprint and the complexity associated with monitoring and managing your storage environment.

Consider the following example: You have a few hundred production VMs on a conventional LUN-based storage platform with custom protection and QoS policies. You need to upgrade your storage hardware and migrate your VMs to the new storage array, requiring you to re-create every LUN with the exact same settings, move the right VMs to the right LUNs, and individually verify that each VM is working correctly on the upgraded platform.

With a VM-aware storage solution, all you need to do is migrate the VMs to the new array and you're finished. This is possible because data protection and QoS policies move with the VMs. VM-aware storage identifies all the VMs in their new home and reapplies the policies.



One of the five key requirements of VM-aware storage (see the spot guide at the end of Chapter 5) is a VM-level filesystem.

So what can you achieve with the ability to move VMs with their policies intact? You can optimize placement of hundreds or thousands of VMs across storage platforms.



Getting the LUN settings or VM placement wrong in conventional storage can negatively affect critical business operations. Storage administrators tend to over-provision or carve out a LUN for the VMs to share, regardless of the storage requirements. You wish you guessed right and didn't have to move them again, which usually doesn't happen.

Regardless of how your business is growing, application requirements are constantly changing (often beyond your

control, such as in the case of hosted workloads). Thus, VMs need the flexibility to move around within your storage array. But you shouldn't have to worry about what underlying storage fits each VM workload best. With VM-aware storage, you can treat all your storage resources as a single datastore pool (see Figure 3-4) and simply move some workloads off an overloaded storage array to another. Or, you can confidently predict that in order to support another 100 virtual desktop infrastructure (VDI) sessions in the next week on the existing array, you need to move the SQL server VMs to another array or add a new array to the pool.

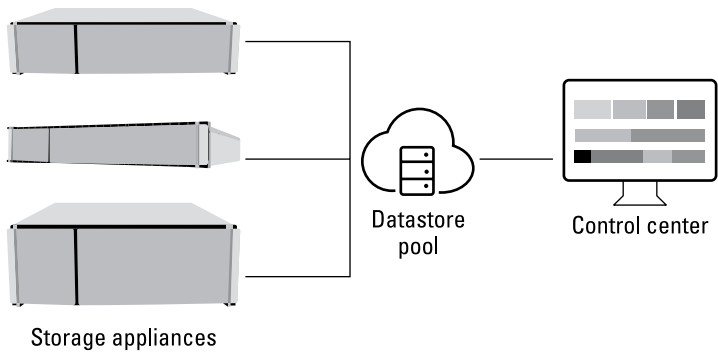



Figure 3-4: Manage workload optimization across datastore pools from a single pane of glass.

In conventional storage, VMs are randomly distributed across different storage platforms, creating performance hotspots and bottlenecks in some storage resources, while other storage resources are under-utilized (see the top illustration in Figure 3-5). This situation is partly caused by “guessing wrong” during the initial planning cycle when assigning VMs and virtualized applications to LUNs or volumes, as well as a lack of visibility into which workloads are growing faster than anticipated once they are deployed.

In contrast, VM-aware storage provides a datastore pool with a clear global view of capacity, performance, and flash usage across all your storage resources in a single management tool. You (or the storage software) can intelligently distribute workloads across the storage infrastructure. Rather than setting an arbitrarily low threshold for moving VMs before a move is really needed, you can make such decisions based on actual VM-level data about growth and usage.

The recommended solution, based on VM storage usage analysis, is to distribute VMs evenly among datastores to avoid hotspots (see the bottom illustration in Figure 3-5). However, if the numerous VMs on the “busy” datastore are small and do not require many storage resources such as performance and flash, VM-aware storage does not blindly move these VMs simply for the sake of even distribution. It makes a smart choice to leave the VMs where they are. This approach ensures that the optimization incurs the least cost in terms of the time and space required to migrate VMs.

Conventional storage



VM-aware storage




Figure 3-5: Comparing management complexity and scalability in conventional and VM-aware storage.

Chapter 4

Exploring VM-Aware Storage Use Cases

.....

In This Chapter

- ▶ Virtualizing business critical applications
 - ▶ Deploying VM-aware storage for the cloud
 - ▶ Taking virtualization to the desktop
 - ▶ Enabling DevOps environments
-

Conventional storage technologies create complexity, cost, and performance issues when used with virtualization. This chapter examines some real-world use cases to help you understand how VM-aware storage addresses these issues.

Hitting Critical Mass with Business-Critical Applications

Before virtualization, business-critical applications, such as customer relationship management (CRM), enterprise resource planning (ERP), and online transaction processing (OLTP)/SAP/Oracle databases were deployed on dedicated hardware, thus creating many isolated islands of hardware that were dedicated to running these applications.

Today, these same applications are being virtualized, and in some cases are being delivered as cloud-based software-as-a-service (SaaS) offerings. With new multi-core hosts and powerful virtualization technology, IT no longer has to choose between performance and the benefits of virtualization. Now they can be achieved together. By virtualizing business critical

applications, you can achieve greater flexibility and reliability with scale up/down compute capacity, distributed resource scheduling, and automated migration and geo-distribution for business continuity and disaster recovery scenarios.

Unfortunately, many of these newly virtualized business-critical applications are attached to conventional storage. Conventional storage doesn't effectively isolate the I/O performance of one VM from another, resulting in blended performance as multiple VMs contend for shared storage resources that were once dedicated resources, before the application was virtualized. Conventional storage is not designed to deal with the unique performance demands of hosting hundreds, or thousands, of VMs.

What these business critical applications require is performance isolation and quality of service (QoS). VM-aware storage is designed to observe the performance profile of the application running in the VM. Using this technology, the VM-aware storage can separate the I/O from each VM, create separate performance queues, and ensure that each VM gets access to the performance and response times that it requires.



One of the five key requirements of VM-aware storage (see the spot guide at the end of Chapter 5) is VM-level QoS.

Microprocessor manufacturer speeds tests and reduces storage footprint

A semiconductor design innovator designs and builds graphics and computing technologies powering a variety of solutions, including PCs, game consoles, and servers.

The conventional storage deployed to support virtualized environments could not meet performance requirements and made scaling for remote sites difficult. Things on the manufacturer's wish list are: a pod-based infrastructure that could be

replicated at remote sites; a flash-based storage system that would satisfy the performance needs of the integrated-circuit testing databases and reduce the storage footprint in their virtualized environments; and visibility into VM performance metrics at the storage layer.

During a storage refresh for the virtualized environments, the manufacturer discovered Tintri. Impressed by the performance in the Tintri

VMstore appliance delivered in a small form factor and its VM-aware functionality, they did a proof-of-concept (POC). Upon successfully concluding the POC, they deployed the Tintri VMstore storage appliance in a pod-based architecture.

Key challenges

- ✔ Insufficient performance made it difficult to run large simulation tests successfully.
- ✔ The company lacked useful metrics for monitoring and troubleshooting the virtual environment.
- ✔ Storage for remote locations required significant space, creating challenges at smaller sites.

Business benefits

- ✔ Ran simulation tests five times faster on Tintri, dramatically reducing completion times.
- ✔ Reduced storage space footprint at the primary site by 67 percent.
- ✔ Saved time because performance and capacity metrics on a per-VM basis allow granular monitoring and troubleshooting.
- ✔ Created pod-based architecture with high VM density for remote locations to reduce operational costs.

VM-aware storage constantly monitors individual VMs. It calculates and automatically assigns the appropriate capacity and performance levels from a pool of underlying storage hardware. VM-aware storage effectively abstracts the complexity of disk management away from the application and simply presents a pool of capacity that system administrators can define, based on required performance metrics set in software.

Enjoying the Benefits of the Cloud (On-premises and Hosted)

Many businesses have transformed parts of their IT operations to the “as a service” model in an on-premises private or hosted cloud environment. This transformation enables organizations to delay or avoid costly investments in specialized infrastructure. In this way, organizations realize the benefits of agility, responsiveness, and scale that virtualized applications can deliver for their business.



A good analogy for comparing traditional IT infrastructure to cloud deployment models is buying and running your own power generator to supply your home's electricity needs. You constantly run your generator at full capacity because you can't store electricity, but you always have to be ready for demand spikes (like when a *Doctor Who* marathon event unexpectedly airs on TV). Or, you could simply buy as much electricity (and enough raw artron energy to power your TARDIS) as you need from a local utility, when you need it.

Conventional storage technologies aren't designed for fast-changing cloud environments, with numerous internal (on-premises private cloud) or external (hosted cloud) customers constantly creating and tearing down applications, and with different resource requirements on a shared storage pool. (Someone inevitably does a cannonball in a crowded pool of small children playing Marco Polo!)

For example, if one customer needs lots of high-performance reads for a database VM while another customer is running a VM with a lot of write operations, such as data backup, conventional storage technologies struggle to cope with these disparate use cases. Performance suffers for all customers.

The management and maintenance overhead of conventional storage can limit the extreme scale required in a cloud environment. Conventional storage offers some automation capability for storage operations and administration, but because it occurs at the LUN or volume level, there is always a risk that automated processes can cause unintended performance consequences in other VMs or applications on the same LUN or volume.

With VM-aware storage, on-premises private and hosted cloud operators can deliver important services to their customers. Quality of services (QoS) can vary between the various VMs to ensure appropriate service-level agreements (SLAs), such as performance and data protection. VM-aware storage allows each VM to be assigned a performance target, which is then automatically delivered through the dynamic use of flash and storage tiering.

Automated QoS policies in VM-aware storage provide each VM with its own I/O "lane." As conditions change, the storage adapts, automatically giving each VM the appropriate level of performance (or width of "lanes"). VM-level backup

and replication technologies are also built directly into the VM-aware storage architecture, which simplifies design and allows for self-service recovery tools.

Compared to conventional storage, VM-aware storage is designed to handle a much larger number of objects (such as VMs, vDisks, snapshots, and the like) and VM-level visibility makes those objects simpler to manage, allowing you to meet cloud-scale requirements far more easily.

Finally, in a VM-aware storage system, automation of administrative tasks, such as provisioning and tearing down virtualized workloads, can be set up at the individual VM level. This capability enables operational efficiencies and massive scalability throughout the organization and for the cloud.



One of the five key requirements of VM-aware storage (see the spot guide at the end of Chapter 5) is VM-level automation.

Cloud service provider chooses “just right” storage solution

A cloud services provider had been leveraging conventional storage arrays with SAS drives for its cloud services, but realized that its existing storage environment would not support the expected growth and required performance for the provider’s new desktop-as-a-service (DaaS) offering.

The company identified several hard requirements, including:

- ✓ Must be able to scale to tens of thousands of desktops and sustain an average desktop load of 40 IOPS at 65 percent write
- ✓ Must be compatible with existing infrastructure and enable

high availability with less than 5 millisecond latencies

- ✓ Must be cost-effective and easy to install and manage

The IT team identified five storage solutions to test, and ultimately selected a Tintri VMstore to replace an all-flash array from the incumbent storage vendor. The Tintri solution provided the best combination of storage capacity, performance, and cost to meet the service provider’s requirements.

Key challenges

- ✓ Conventional storage required a large upfront investment in expensive controllers.

(continued)

(continued)

- ✔ Deploying legacy storage arrays to production took at least two to three days for configuring, racking, installing power and cabling, and other deployment tasks, and required specialized administrator training.
- ✔ Delivered significant performance improvements to end-users
- ✔ Provided linear scalability of capacity into production in 85 percent less time than previous conventional storage
- Business benefits**
- ✔ Simplified storage management and eliminated the need for specialized training
- ✔ Obtained VM-aware storage at a 35 percent lower cost per desktop

Empowering Your Workforce with VDI (That Actually Works)

A *virtual desktop infrastructure* (VDI) replaces traditional desktop PCs with a virtual desktop interface that is accessed across a network (or Internet) connection, and actually runs as a VM in the data center.

The advantages of VDI are greatly amplified when a large number of end-users require desktop access, such as in colleges, hospitals, libraries, and call centers, or in situations where users are geographically dispersed (for example, home offices or remote branch offices with little or no local IT support).



The maintenance of hardware, software, security, and data management in a VDI environment can all be centralized for more efficient and secure administration. For example, security updates can be installed on a single VDI image that is used by everyone, instead of being installed individually on hundreds or thousands of desktop PCs.

With VDI, users can securely access their virtual desktop and applications from a variety of endpoints, such as laptop or desktop PCs, tablets, and smartphones. By enabling greater mobility and enterprise “bring your own device” (BYOD) policies, VDI helps to promote flexible workplace practices and drive greater productivity.



VDI places extreme demands on conventional storage, including:

- **Multiple concurrent logins:** When employees log on to their virtual desktops to begin their workday, the high volume of simultaneous logins can cause a *boot/login storm* resulting in long delays.
- **Dynamic workloads:** Unlike other common workloads that often follow predictable patterns, each VDI user accessing applications and data may have varying input/output operations per second (IOPS) requirements and block size patterns, which makes tuning conventional storage difficult.
- **Over-provisioning:** To meet performance demands, VDI deployments using conventional storage technologies must allocate more storage resources than are typically required to overcome boot/login storms and unpredictable workloads. Over-provisioning is a wasteful and costly strategy.



A VM-aware storage layer is well suited to VDI deployments in several ways:

- **Flash performance:** Flash storage technology substantially lowers the performance impact of boot and login storms while providing quick and consistent response times for dynamic VDI workloads.
- **Deduplication and compression technology:** This feature removes repeated content and shrinks the amount of data stored, greatly enhancing the storage efficiency.
- **Full visibility at the VM level (each VDI instance):** Administrators can accurately determine whether performance thresholds and SLAs are being met, while finding and throttling noisy neighbors that may be adversely affecting the VDI environment. This visibility also allows the VM-aware storage system to host workloads with differing I/O performance characteristics and to dynamically allocate resources between them as required among workloads.



The combination of VM-aware flash storage and data optimization techniques provides a huge performance boost that negates the issues caused by highly dynamic VDI workloads while achieving high storage efficiency.

University graduates to VM-aware storage for campus-wide VDI deployment

An independent, private university with a very small technology budget and IT staff supports 350 faculty and staff, 1,000 undergraduate students, and up to 3,000 graduate students at the main campus and two satellite campus locations. Although the first phase of the virtual desktop project was successful, performance was an issue for the persistent desktop deployment.

The team decided to test six different storage platforms to find the best solution for the VMware desktops. In addition to conventional SAN storage, they looked at pure flash offerings, but none of those offerings were feasible with the university's limited IT budget. The Tintri solution outperformed all other vendors' products in the POC, leading the university to purchase a single Tintri VMstore for its VMware VDI environment.

Key challenges

- ✓ Latency in the existing VDI environment

- ✓ Conventional storage could not provide sufficient IOPS
- ✓ High maintenance costs associated with desktop PCs
- ✓ Desktop PCs required too much power

Business benefits

- ✓ Eliminated boot storms and latency issues in the VDI environment with 100 percent flash hits
- ✓ Cut time to update applications from 50 hours to 10 minutes
- ✓ Provided anywhere, anytime desktop access in support of the university's "bring your own device" (BYOD) initiative
- ✓ Reduced power consumption from 120 watts for each traditional desktop PC to 3 to 5 watts per virtual desktop — a 95 percent reduction
- ✓ Saved CapEx and OpEx

Benefiting from DevOps with Agile and Grace

Companies seeking to gain a competitive edge or to solve challenging problems with technology often devote significant engineering resources to developing and testing application software. The speed with which these applications can be

developed and tested greatly determines the productivity of the engineering team and the value of the application software. An application that is late to market has no chance of creating a competitive advantage.

Development teams increase their productivity by rapidly prototyping and testing new application software against a wide range of deployment scenarios, production workloads, and potential failure modes. Having large numbers of test and development environments available to development teams can significantly reduce the time required to deliver new applications, but require significant capital expenditure, setup, and resource contention management.

Virtualization enables rapid prototyping and large-scale simultaneous testing in support of multiple concurrent projects. Unfortunately, conventional storage platforms don't have the flexibility needed to build virtual test and development environments. Although fine for a single static application, conventional storage suffers from scalability issues, performance bottlenecks, and management complexity when used to support the dynamic creation and teardown of thousands of VMs in large-scale test and development environments. Conventional storage requires extensive setup, configuration, and maintenance work by software, virtualization, and storage teams, effectively limiting the organization's ability to conduct rapid prototyping and testing.



One of the five key requirements of VM-aware storage (see the spot guide at the end of Chapter 5) is VM-level storage operations.

A complete DevOps solution combines server virtualization and VM-aware storage. A VM-aware storage solution is designed to support VMs as objects, with features that are particularly well-suited to test and development environments, including:

- ✓ **Powerful tools to clone VMs:** Server virtualization workflows that create thousands of copies of development and test VMs can be offloaded to automated storage cloning tools, requiring seconds to deploy, rather than hours or days.
- ✓ **Self-provisioning storage to support DevOps workflows:** The rapid deployment of test and development VMs is possible only when storage administrators aren't

required to manually configure storage or decide where to place new VMs. With VM-aware storage, storage administrators can create a pool of storage for developers to consume as required and then release when development projects scale down and the VMs are no longer needed.

- ✔ **Robust tools to refresh development VMs:** Refreshing a test and development system should not require the entire development environment to be re-created. VM-aware storage offers tools that can revert the data to a previous snapshot in a VM without replacing or rebuilding the VM. With the correct VM-aware tool, this update process takes seconds, not hours.

EU bank selects Tintri for its virtualized test and development environment

A European Union (EU) bank launched its enterprise-wide virtualization project beginning with test and development systems. The physical environment was running on conventional storage that was not capable of meeting the increased performance demands of a virtualized environment.

After looking at different storage options, the bank purchased six Tintri VMstores and Tintri Global Center for its virtual test and development environment.

Key challenges

- ✔ The bank needed a high-performance storage platform for testing and development.
- ✔ Existing conventional storage was not up to the demands of virtualization.

- ✔ The total cost of ownership (TCO) for conventional storage appliances in the production environment was high.

Business benefits

- ✔ Achieved storage performance for testing and development that was comparable to production systems.
- ✔ Simplified storage management and troubleshooting with full visibility into performance of individual VMs and zero downtime associated with upgrades.
- ✔ Reduced power consumption in each data center by 5 percent.

Chapter 5

Ten Cool Things You Can Do Only with VM-Aware Storage

In This Chapter

- ▶ Boosting performance and productivity
- ▶ Getting “just the facts, ma’am” with analytics
- ▶ Reducing cost and complexity

Congratulations! You’ve reached the end of *VM-Aware Storage For Dummies* (or, perhaps you started with this chapter). In either case, in this chapter I present . . . drum roll please . . . ten cool things you can only do with VM-aware storage in that classic “For Dummies” style.

- ✔ **Get rid of noisy neighbors (guaranteed performance).** VMs stuffed into a conventional storage logical unit number (LUN) share the same performance characteristics and policies assigned to that entire LUN. When one of those VMs overly demanding or goes rogue, you’ve got a noisy neighbor. VM-aware storage eliminates LUNs, instead assigning every VM its own lane, which includes performance isolation. VM-level quality of service (QoS), helps set the right values for the width of those lanes, so every VM gets the appropriate speed min or max.
- ✔ **Perform surgery with a scalpel instead of a chainsaw (end-to-end visibility).** VM-aware storage enables you to manage every VM as an individual object and shows latency across hosts, network, and storage. Never

confuse symptoms for the root cause — the end-to-end visibility of VM-aware storage lets you visualize contention across infrastructure and solve problems in a few clicks.

✔ **Click three times to send your VMs back to Kansas (per-VM replication).** Per-VM replication is a simple but powerful concept. Instead of replicating an entire LUN (containing hundreds of VMs) in conventional storage, you can replicate only the individual VMs you need with VM-aware storage, which is what you should be doing. Now close your eyes, imagine replicating across multiple sites to a remote disaster recovery location, say “there’s no place like home” and click your heels (er, mouse) three times to achieve:

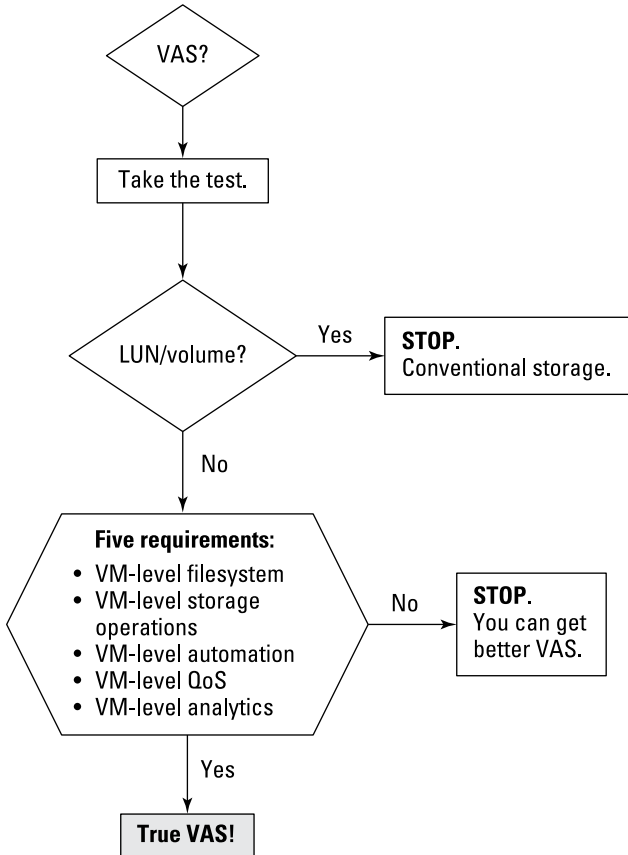
- Recovery-point-objectives (RPOs) as low as one minute for your mission-critical applications
- WAN bandwidth reduction of up to 95 percent with integrated deduplication and compression

✔ **Travel through time without a flux capacitor (efficient copy data management).** With conventional storage, application development teams contend with slow and inefficient processes when trying to refresh their VM-based environments with new production data because it requires copying data on the entire LUN, identifying target VMs within the snapshots, and reconfiguring existing VMs to use new data sets. Buy your developers a DeLorean and they’ll love you forever. Okay, maybe buy them VM-aware storage instead, so they can instantly recover any point-in-time version of a VM and preserve snapshot and performance history of the VM between versions.

✔ **Enjoy clones whether or not you’re a “true” *Star Wars* fan (VM-level data management across hypervisors).** With VM-aware storage, per-VM snapshots and clones can be scheduled automatically by default — with no performance impact. Because all data services are per-VM, there is no spillover to other VMs. Instant provisioning and space-efficient clones are available to hypervisors nearly instantly. Jedi administrators can create on-demand fully functional clone VMs to remote VMstore systems more efficiently.

- ✔ **Get the temperature for your local city rather than the entire continent (VM-level analytics).** If someone claims to deliver VM-level analytics, check under the covers. It usually only takes some basic math (the stuff you learned in third grade) to figure out what is used: averages of latency, throughput, and other storage performance characteristics across all the VMs in a LUN (translation: “It’s still LUN-level analytics!”). True VM-level analytics allow you to dig into the performance characteristics of an individual VM with precision, not estimations or correlations.
- ✔ **Maintain perfect balance without yoga (VM scale-out).** With VM-aware storage, you can scale out storage performance and capacity, knowing that the unique requirements of each individual VM or application have been met and optimized. So save your awesome Lotus Pose for marathon budget meetings!
- ✔ **Save enough money to buy a small island (no more over-provisioning).** I’m talking about an island with beaches and sand, not an island of storage with silicon! With VM-aware storage you can better forecast your storage requirements and accurately plan your capital investments, without costly and wasteful over-provisioning.
- ✔ **Don’t let them find Sasquatch in your data center (go “green”).** When it comes to your carbon footprint, you don’t want to be Bigfoot. If you haven’t Yeti taken virtualization to the next level, you can achieve unparalleled VM density by consolidating thousands of virtualized workloads in as little as two rack units (RUs) on VM-aware storage, which saves power, cooling, and real estate OPEX — as well as the environment.
- ✔ **Spend your weekends doing what you love — as long as it isn’t managing storage (eliminate management complexity).** If you enjoy carving out LUNs and whiteboarding storage performance issues, it might be time to find a new hobby. VM-aware storage restores the logical 1:1 mapping between virtualized applications and storage, and simplifies storage management in virtualized and cloud environments at scale.

VM-Aware Storage (VAS) Spot Guide



S E E S T O R A G E D I F F E R E N T L Y



**Tintri VM-aware storage is the simplest
for virtualized applications and cloud.**

Organizations including Toyota, United Healthcare, NASA and 7 of the Fortune 15 have said “No to LUNs.” With Tintri they manage only virtual machines, in a fraction of the footprint and at far lower cost than conventional storage. For more information, visit www.tintri.com and follow us on Twitter: @tintri



Create a successful storage architecture for supporting virtualized apps


There's a disconnect in your data center. Your physical-era storage isn't built to effectively support the virtualized applications that matter to your business. Don't despair — a new generation of VM-aware storage is built specifically for virtualization and cloud.

Read all about it in this book.

- *Seek innovative alternatives to convention* — address the disconnect between virtualized applications and conventional storage
- *Discover VM-aware storage* — focus on the applications that matter most for your business
- *Maximize the value of VM-aware storage* — witness how storage can solve your headaches, and discover business benefits far beyond
- *Explore real-world use cases* — learn how other businesses have put VM-aware storage to work
- *“Wow” your colleagues* — try ten cool things you can do only with VM-aware storage

Lawrence C. Miller has worked in information technology for more than 25 years. He has written more than 50 For Dummies books.

FOR
DUMMIES[®]
A Wiley Brand

 Also available
as an e-book

ISBN: 978-1-119-28856-5
Not for resale



Open the book and find:

- Differences between conventional storage and VM-aware storage
- Best practice tips for VM-aware storage
- Practical examples of VM-aware storage in the real world
- What you need to know to build VM-aware storage for virtualization and cloud

Go to **Dummies.com**[®]
for more!

WILEY END USER LICENSE AGREEMENT

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