

Making Everything Easier!™

Tintri Special Edition

Application-Aware Storage

FOR
DUMMIES®
A Wiley Brand

Brought to you by



Tintri

**Saradhi Sreegiriraju
Vineet Kakani
Will Garside**



About Tintri

A new model for IT is here. Virtualized applications are the norm, and traditional storage is 20 years overdue for a shakeup. Tintri builds smart storage that sees, learns, and adapts, enabling IT organizations to focus on virtualized applications and business services instead of managing storage infrastructure. Tintri application-aware storage eliminates planning and complex troubleshooting by providing VM-level visibility, control, insight, and agility, moving IT from reactive to strategic. Tintri improves performance by 10x, lowers OPEX by 60x, and CAPEX by 10x, and frees IT professionals from the drudgery and complexity of traditional storage. Tintri powers hundreds of thousands of virtual machines running business-critical databases, enterprise apps, desktops, mobile apps, and private cloud deployments. Tintri helps global enterprises such as AMD, F5 Networks, GE, NEC, NTT, MillerCoors, and Time Warner maximize their virtualization and cloud investments.

For more information, visit www.tintri.com and follow us on Twitter: @tintriinc or @Tintri_EMEA.

Application-Aware Storage

FOR
DUMMIES®
A Wiley Brand

**by Saradhi Sreegiriraju,
Vineet Kakani and Will Garside**

FOR
DUMMIES®
A Wiley Brand

Application-Aware Storage For Dummies®

Published by
John Wiley & Sons, Ltd
The Atrium
Southern Gate
Chichester
West Sussex
PO19 8SQ
England

For details on how to create a custom *For Dummies* book for your business or organisation, contact CorporateDevelopment@wiley.com. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

Visit our Home Page on www.customdummies.com

Copyright © 2014 by John Wiley & Sons Ltd, Chichester, West Sussex, England

All Rights Reserved. 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 under the terms of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency Ltd, 90 Tottenham Court Road, London, W1T 4LP, UK, without the permission in writing of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Ltd, The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, England, or emailed to permreq@wiley.com, or faxed to (44) 1243 770620.

Trademarks: Wiley, the Wiley Publishing logo, For Dummies, the Dummies Man logo, A Reference for the Rest of Us!, The Dummies Way, Dummies Daily, The Fun and Easy Way, Dummies.com 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. Wiley Publishing, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER, THE AUTHOR, AND ANYONE ELSE INVOLVED IN PREPARING THIS WORK 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 FROM 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.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

ISBN: 978-1-118-98521-2 (pbk); 978-1-118-98520-5 (ebk)

Printed and bound in Great Britain by Page Bros, Norwich

10 9 8 7 6 5 4 3 2 1

Introduction



Welcome to *Application-Aware Storage For Dummies* – your guide to understanding the key concepts and technologies involved in creating a successful storage architecture for supporting virtualized applications and workflows.

About This Book

This book may be small, but we've packed it full of helpful guidance on how you can design, implement, and manage valuable data and storage platforms.

Foolish Assumptions

In writing this book, we've made some assumptions about you. We assume that:

- ✓ You're a member of an organization planning to implement a virtualization and data storage project.
- ✓ You may be a manager or team member but you aren't necessarily a technical expert.
- ✓ You need to be able to get involved in the project in some way, and you may have a critical role in which you can benefit from a broad understanding of the key concepts.

How This Book is Organized

Application-Aware Storage For Dummies is divided into five concise and information-packed chapters:

- ✓ **Chapter 1: Identifying the Challenges of Managing Traditional Storage within Modern IT Environments.** This chapter walks you through the fundamentals of

data storage technologies and common problems you can encounter.

- ✓ **Chapter 2: Solving Traditional Storage Issues with Application Awareness.** This chapter helps you to understand how to solve five big legacy storage challenges whilst delivering additional benefits, too.
- ✓ **Chapter 3: Building Effective Virtualized Infrastructures Using Application-Aware Storage.** Find out in this chapter how Application-Aware Storage can solve common use problem cases, and get the lowdown on some real-world examples.
- ✓ **Chapter 4: Understanding Key Management Concepts to Get the Most Out of Application-Aware Storage.** This chapter serves up advice for keeping your project on track and for providing long-term management best practice.
- ✓ **Chapter 5: Ten Tips for a Successful Migration to Application-Aware Storage.** The top tips offered in this chapter zoom in on the areas that you need to remember in order to ensure success.

Icons Used in This Book

To make it even easier for you to navigate to the most useful information in this book, these icons highlight key text:



The target draws your attention to top-notch advice.



The knotted string highlights important information to bear in mind.



Watch out for these potential pitfalls.

Where to Go from Here

You can take the traditional route and read this book straight through from cover to cover. Alternatively, you can skip between sections as you like, using the headings as your guides for pinpointing the information you need.

Whichever way you choose, you can't go wrong. Both paths lead to the same outcome – the knowledge you need to build a highly scalable, easily managed, and well-protected Application-Aware Storage architecture for supporting critical applications and workflows in a virtualized environment.

Chapter 1

Identifying the Challenges of Managing Traditional Storage within Modern IT Environments

In This Chapter

- ▶ Appreciating how applications have benefited from server virtualization
 - ▶ Grasping the problems with traditional storage in virtualized environments
-

Modern organizations need electronic systems to manage increasingly complex processes. Application software running on computer servers with access to data storage is one such process and a vital resource for uses such as order processing, inventory control, payroll, and a host of other tasks.

Virtualization – an abstraction where the resources of a single physical computing server are split into multiple independent virtual servers – has led to more agile server technology designs for hosting application software being introduced, to the extent that *virtual machines* (or VMs – software computers that, like physical computers, run an operating system and a business application) have become proxies for the applications hosted within them. However, the data storage technology for running these virtualized applications hasn't kept up at quite the same pace.

Speaking the IT lingo

Before you dive headlong into the details, getting to grips with basic IT terminology is a 'must':

- ✓ **Data** is a term for formatted information collected, organized, and often processed by a business to help perform a useful task.
- ✓ **Software applications** are able to carry out complex mathematical and process-driven tasks, often using collected data to help meet business drivers.
- ✓ Software applications run on a technology layer of physical hardware, such as powerful computing systems called **servers** and dedicated data storage hardware.
- ✓ Linking applications, servers, and storage devices is a **communication network** that enables the delivery of enterprise processes.

In this chapter, we get you up to speed with virtualization and the challenges that traditional storage methods present within a virtualized environment.

Appreciating the Impact of Virtualization

Only until relatively recently, software applications typically ran as a single instance on physical server hardware. These software applications – such as databases, for example – stored data on a separate physical storage device and sent data items across a connecting network as needed.

Understanding the basic limitations of physical server deployments



Deploying an application on its own physical server hardware has the following limitations:

- ✓ **Poor scaling of servers to meet growing demand.** The traditional method of having a single application on underutilized server hardware makes it difficult to scale the application to meet increased demands or more complex tasks. Increasing the performance of an application would require more powerful underlying hardware. Alternatively, additional individual application servers could be installed and then connected together in close proximity to form a computer cluster, but this is a more technically complex process to create and manage.
- ✓ **Underutilization of server hardware.** Server hardware has become incredibly powerful over the last few years but the software running on this hardware is often unable to fully utilize all the available processing power. This underutilization of server hardware is expensive and wasteful in terms of space and power consumption.

The revolution will be virtualized

Back in the 1990s, a new technology called *virtualization* revolutionized the traditional computing model (see Figure 1-1). Virtualization introduced clever software, called *hypervisor*, which would effectively split up the resources of a single physical computing server into multiple virtual servers. Known as *virtual machines* (VM), these virtual servers can reside within the same physical hardware and each be assigned a variable amount of processing capacity to match the demands of the applications.



Virtualization greatly improved server resource utilization. Whereas previously, each software application typically would run on independent physical server hardware, virtualization meant that multiple virtual machines could now run multiple virtualized applications on one physical server, better utilizing the available compute capacity. Virtual machines could be created quickly and be moved around between physical servers, allowing organizations greater flexibility in where and how they scaled and processed their workloads.




Figure 1-1: Physical server architecture versus virtualized computing architecture.

Understanding Problems with Traditional Storage in Virtualized Environments

The rise of server virtualization has dramatically improved the efficiency of many IT processes, including data storage. Whereas traditional storage methods map one-to-one between applications and storage, virtualization enables dozens of servers to share a common storage pool. Traditional methods used for data storage weren't designed with virtualization in mind and have proved to be a force fit, at best, for the highly virtualized data centers.

Meeting the different types of traditional storage

Traditional storage architectures can be divided into three main types (see also Figure 1-2):




Figure 1-2: Comparing the three main types of traditional data storage.



- ✓ **Direct Attached Storage (DAS)** is storage capacity located either inside the same physical server hardware or directly attached via a cable that doesn't pass via a general-purpose data network connection. DAS is common in smaller-application servers and for working with very small data sets during simpler tasks, and mandates no sharing of storage or data with other application servers.

DAS is potentially more vulnerable to data corruption and loss than network connected storage and because DAS is reliant on the available resources of its host server it has only limited ability to scale up in performance.

- ✓ **Network Attached Storage (NAS)** is data storage capacity connected to a network in order to allow data exchange between one or more application servers and clients. The NAS consists of multiple data storage disks holding duplicate copies of the data to protect against data loss in the event of one or more of the disks suffering a mechanical failure. The majority of NAS devices



use *Ethernet*, the most popular type of general-purpose network, to send and receive data between application servers and users.

Each NAS appliance is effectively a standalone device. When the NAS reaches its maximum capacity or performance limit, upgrading to a bigger NAS or adding additional NAS appliances to the infrastructure can add further resource. These multiple NAS devices must then be managed and applications reconfigured in order to make sense of how and where the data now resides.

- ✓ **Storage Area Networks (SAN)** are similar to NAS, but use a dedicated network normally based on optical fiber to create a joined-up pool of storage.



SAN technology potentially has greater storage capacity and performance than NAS, but at a higher overall cost and with slightly more complexity.

Identifying problems with traditional storage within virtualized environments

DAS-, NAS-, and SAN-based traditional storage technologies were created in the era before virtualization and weren't designed to support the characteristics of a virtualized workload. All three have limitations that restrict their effective use within hybrid and predominantly virtualized environments.



These problems include:

- ✓ **Complex management.** Virtualization is effectively managed through software tools that enable organizations to create new VMs, move workloads, and – increasingly – assign virtualized networking resources. However, traditional storage management tools, built before the rise of virtualization, are not well suited to management

processes that have become commonplace within virtualized environments. For example, in some instances, virtualization enables workloads to scale up and down automatically, based on demand, yet most traditional storage technologies require manual intervention to assign the corresponding storage capacity before the workload can start.

✓ **Inflexibility in assigning resources.** One of the major advantages of virtualized environments is the ability to quickly adapt to new requirements. Different types of use cases require data to be stored and made available in a variety of tiers. These *tiers* describe an access pattern built from a number of different storage criteria, which can include:

- Ability to perform data reduction
- Accessibility for multiple users
- Cost per unit of capacity
- Length of delay before data becomes available for use
- Resiliency against data corruption
- Read performance
- Write performance

Traditional storage mandates a number of complex steps to reconfigure data storage and services to meet tiering characteristics of dynamically changing needs of virtualized environments.

✓ **Lack of ability to scale capacity as needed.** The use of virtualization and the ease of server provisioning has led to 'VM sprawl', with hundreds of instances of applications being easily created 'on the fly'. However, these applications are still dependent on data from traditional storage. To configure traditional storage to the needs of an application running on virtualized servers requires cumbersome, lengthy, and time-consuming manual

steps. In addition, unless wasteful over-provisioning has occurred, the traditional storage may not have the right performance characteristics to support hundreds of virtualized servers hosting hundreds of applications.

✓ **Poor visibility of performance characteristics.** In the era before server virtualization, traditional storage infrastructure was effectively mapped against single monolithic applications on single physical application servers. However, in a virtualized environment with multiple VMs, each server houses multiple applications. This setup creates the *IO blender effect* (IO meaning ‘input/output’) in which operations from applications housed in multiple VMs are mapped to one storage container. This makes understanding what each software application needs, in terms of data storage capacity and performance, more difficult. This lack of visibility between the application and the storage infrastructure makes finding and solving any performance issues more difficult. (Figure 1-3 illustrates the IO blender effect in action.)




Figure 1-3: Comparing application deployments in physical and virtualized environments on traditional storage.

✓ **Waste due to over-provisioning.** Virtualized environments tend to evolve quickly. However, to meet performance requirements and application infrastructure needs, IT system designers have to buy more storage hardware with more storage capacity than they actually need – a process called *over-provisioning*. Over-provisioning is incredibly wasteful in terms of cost and space as much of the capacity purchased may not be needed for months or even years, but has to be bought ‘up front’.

Chapter 2

Solving Traditional Storage Issues with Application Awareness

In This Chapter

- ▶ Understanding application-aware storage
 - ▶ Grasping the changes that virtualization has brought to storage
 - ▶ Confronting the challenges of traditional storage with application-aware storage
-

In recent years, a new type of storage technology – *application-aware storage*, designed to meet the needs of a wide range of uses supported by modern computing systems – has been addressing the mismatch between server virtualization and traditional data storage technology. Application-aware storage can help to solve the pain points suffered by organizations who are trying to benefit from switching to a virtualized environment. (You can read about these pain points in Chapter 1.)

This chapter looks at how these organizations can benefit, and serves up a simple explanation of the basic elements that make up the architectures of application-aware storage.

Getting to Grips with Application-Aware Storage

In the context of data storage, application-aware storage is a system that has built-in intelligence about the relevant applications that wish to use it. This intelligence can include a number of factors, such as an understanding of:

- ✓ Acceptable delays in finding data (called *latency*)
- ✓ Data availability
- ✓ Data reading patterns
- ✓ Data writing patterns
- ✓ How to protect the application in various failure scenarios



When the storage system ‘understands’ the applications and usage conditions, it can optimize how data is stored and use additional technologies, such as deduplication and compression, to improve performance or quality of service levels to meet the needs of different use cases. Such understanding also enables you to independently manage and optimize the storage for each individual application, regardless of what other workloads and applications are present within the wider IT environment.

Table 2-1 gives a few examples of the varying storage requirements that you might have in different application use cases. However, the list can also include a host of other potential elements such as:

- ✓ 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

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

	<i>Read perform- ance</i>	<i>Write perform- ance</i>	<i>Random access perform- ance</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

Application-aware storage is able to work in conjunction with the virtualization infrastructure to determine and understand what individual virtualized applications demand from the storage system. This insight is then used to make the storage system more responsive to the demands of those applications.



To give an analogy from everyday life, try to imagine that you had the ability to examine all the vehicles that were travelling through a particular town every day. If you could understand where they were all starting from, going to, and what the purpose of their journeys was, then you could organize the traffic light patterns across town in a way that would enable all those journeys to flow more smoothly. Furthermore, if you had the additional ability to change which lanes were open or to turn single lanes into dual carriageways running in a particular direction during a certain time of day, you would eliminate traffic jams. Just imagine! Your understanding of the traffic patterns would enable you to better optimize the roadways. In a virtualized environment, this is the kind of power that application-aware storage provides.

Understanding How Virtualization has Changed Storage Requirements

In the era before application virtualization, tuning storage to meet the unique constraints of applications was a complex, but understood, process. Generally speaking, each application had dedicated storage hardware with capacity and performance characteristics designed to meet the application's requirements.

Fast-forward to the era of virtualization, with multiple applications running in multiple virtualized machines (VMs), all of which are sharing storage, and you find that it isn't possible to specifically allocate dedicated storage hardware to each application to meet the required capacity and performance characteristics. An initial compromise was to define a small number of fixed service levels that each virtual machine could be assigned, (as shown in Figure 2-1) but inherent difficulties in understanding whether a virtual machine needed its assigned level or was fully utilizing the committed resources made this a wasteful technique.




Figure 2-1: Virtual machines with various service levels.

What was needed was *application-aware storage* to cater to the needs of applications running inside the VMs. Enter *virtual disks*, which VMs use for their operating system, application software, and other data files. A virtual disk hides the physical storage system from the VM's operating system. Regardless of the type of storage device that the host uses, the virtual disk always appears to the VM as a local storage device. As a result, the VM can run operating systems that are not certified for specific storage equipment in the VM. This virtualization of the storage system improves flexibility by enabling organizations to use multiple types of storage technologies within the same virtualized infrastructure to take advantage of different characteristics, such as lower per-unit capacity costs or better performance or reliability attributes.



Associated storage management tasks used to be difficult to understand because of the separation of storage container from the host. Application-aware storage eliminates this issue by acting as a dedicated resource manager, delivering different levels of performance and quality of service depending on the needs of the application and broader requirements of the virtualized IT environment.

Solving the Five Big Challenges of Traditional Storage with Application-Aware Storage

If you've read Chapter 1, cast your mind back and you'll recall finding out about the five major issues with traditional storage within virtualized environments. Application-aware storage enables you to overcome those challenges in several ways.

Simplifying management complexity

The complexity of storage management stems from the lack of suitability to the task that traditional storage offers in an increasingly virtualized environment. In the traditional application and storage world, a manual process defined how you managed data (called the *control plane*) and where that data

physically resided (called the *data plane*). You had to assign to each application a certain storage capacity and physical path, and both adding and removing resources required many administrative steps. In order for them to function effectively, you had to precisely tell the operating systems that underpinned applications – and even the applications themselves – the characteristics of the underlying storage.

In many traditional storage architectures, the control plane was effectively wired into the physical elements of the hardware. Making fundamental changes would require wiring up the hardware in a different manner. Recent years have witnessed a rise in placing more of this control intelligence within software, combined with advanced communication methods in order to abstract many of these complex manual processes into software-based alternatives. In these *software-defined architectures*, in effect, the control plane is now separate from the data plane, which helps to remove many of these restrictive and complex management issues.



Application-aware storage uses many of the principles of software-defined architecture, but takes them a step further. Many of the software tools for managing the storage system's control plane can be assisted by information gleaned from interrogating the VMs. In some instances, the software-defined controls may be automated based on the VM reporting certain threshold values (such as performance, available capacity, and latency) that might trigger the automated provisioning or make more resources available.

Introducing flexibility in assigning resources

The software-defined flexibility offered by application-aware storage helps to solve the problem with setting fixed tiers of storage service that persisted in traditional storage. A deeper understanding of what an application demands of the storage enables application-aware storage to dynamically alter the storage environment to meet requirements of the application.



The ability to deliver a particular quality of service can be specified down to the individual application, and can be continually updated based on a number of conditions. For example:

- ✓ **A particularly critical web application is demanding more performance.** The application-aware storage looks at what else is running. Noticing that a non-essential data archive job could be allocated fewer resources, it diverts the freed-up surplus resources and helps to serve data to the demanding web application.
- ✓ **A backup job needs to complete during an overnight period.** The application-aware storage may reduce the resources available to all other applications by a small percentage so as to divert more resources to the crucial overnight task.
- ✓ **In the morning, school students turn on their PCs, creating a sudden demand for a number of critical files to be available from storage as part of the start-up process.** The application-aware storage keeps a copy of these files in its flash cache each morning between 9 a.m. and 10 a.m., but after this point it releases these files to prioritize other workloads.

Tackling a lack of ability to scale capacity as needed

Modern enterprises are facing growing data storage requirements. In part, this situation is due to an increased reliance on digital processing and communication tools and greater data retention for business intelligence projects. However, accurately predicting how much data storage capacity each unique organization will need in five years, let alone in one, is incredibly difficult, and the ability to scale capacity is constrained by factors such as available space, power, and cooling.

Another capacity consideration is that storage technology is advancing very quickly. The house brick-sized hard disk storage devices of the 1980s, holding the data equivalent of a single encyclopedia, have been replaced by devices an eighth of that size which are able to store the text from every book ever written!

Server virtualization helped to solve one of the major challenges of scaling applications by getting more usage out of available computing power. To solve scaling challenges with

traditional storage, however, the storage industry is starting to embrace scale-out techniques which can also benefit from application-aware storage.

A *scale-out architecture* assumes that you want to grow capacity over time and that each additional upgrade you make will be add both storage capacity and performance capability. Scale-out uses a type of building brick methodology that adds each appliance into a seamless namespace so that the underlying virtual servers and applications can access more capacity but without requiring reconfigurations.



The addition of application awareness to this scale-out architecture enables the upgrade process to be more granular. This is because application awareness gives an organization better insight into the underlying workloads and demands of applications compared to their ability to deliver against the requirements offered by the storage system. So, you might configure the type of storage appliances added to storage pools for more flash or for slower, but cheaper, hard disks for workloads where performance isn't needed. The scale-out may favour appliances with more connectivity to the network, to reduce issues like latency, or to support more concurrent users. This mix-and-match approach enables organisations to tailor their requirements and budgets with more finesse.

Seeing off poor visibility of performance characteristics

The IO blender effect (see Chapter 1) creates a lack of visibility into individual applications running as VMs across traditional storage systems. This makes troubleshooting very time consuming for administrators who are forced to investigate problems using multiple and separate tools for applications and storage systems.

In large virtualized environments, troubleshooting performance issues is very much a process of trial and error, because the storage doesn't know anything about the application – it is effectively 'dumb'. Application-aware storage, however, is given a template built around the requirements of the applications that use it, and it then shapes a set of storage resources (such as flash, compression, deduplication, and backup and recovery procedures) to

meet this template. Application-aware storage acts as a storage container that tracks all the characteristics of resident VMs and can correlate with the rest of the virtualized infrastructure in real time. Hence, it can help to build the usage template and validate that the storage is meeting the demands the virtualized application has placed upon it. Since the application-aware storage gathers and presents consolidated information at the individual application level, administrators are better able to spot issues and enact effective solutions in much less time than with traditional storage systems. Figure 2-2 illustrates the mapping between applications and storage in physical environments, virtualized environments using traditional storage, and virtualized environments using application-aware storage.




Figure 2-2: Comparing application deployments in physical and virtualized environments using traditional and application-aware storage.

For example, take firstly the scenario of an application starting to run slowly. The application-aware storage could present data indicating that the VM is experiencing a delay in data being delivered to the applications inside it because of issues in the network layer. Also consider the example of an application-aware storage indicating that a particular VM is misconfigured, meaning that the application running inside the VM is experiencing slow performance. An interrogation

of the VM highlights that, in fact, this particular machine was not configured with the proper amount of computer power for the application running inside and that a reconfiguration is required to fix the issue.

Getting rid of waste caused by over-provisioning

There's no doubt about it – trying to match up different storage characteristics such as performance, capacity, and cost to application requirements is a tricky task. Getting it wrong simply results in wasted resources. However, application-aware storage intelligently utilizes the following three technologies to reduce over-provisioning, to improve performance and to shrink the volume of data that the systems need to retain to improve efficiency:

✓ **Flash storage** uses memory chips or solid-state drives instead of slower spinning magnetic disks to store information. As such, you can more quickly store and retrieve data to improve application performance. However, a downside is that flash storage is many times more expensive, per unit of capacity, than disk drives.

In an application-aware storage environment, flash is used as a type of fast storage area for frequently used information. The application-aware storage can intelligently prioritize what data should be stored within the flash portion of a storage appliance and thereby help to boost performance, based on the demands of the application. So, for example, say that a company has a critical stock file that is often accessed and amended by a stock control system. Instead of this file sitting on slower disk storage, this file might live permanently on flash in order to improve performance.

✓ **Deduplication** spots identical stored information, removes these duplicates, and leaves only a pointer saying that a copy of the duplicate data is stored at a particular location within the storage architecture. Say, for example, that a memo gets circulated around a company via email and lots of staff copy it onto their local computer systems. However, every week these local systems are backed up and potentially hundreds of copies of this

same memo will be stored onto the backup storage systems. Deduplication would notice that one copy of this memo is already on the system. With any subsequent attempts to back up the memo, deduplication technology would instead leave a little pointer to the computer with directions of how to automatically retrieve the source copy. Deduplication can also take place as files are copied into the disk or at a later stage.

✓ **Compression** is a mathematical technique that looks at the structure of data to find patterns that can be better described using mathematical coding. For example, a typical word-processed document might include the 500 most popular English words within 40 per cent of the stored document. These multiple character words could instead be represented as a single special character instead of the 3-, 4-, 5- and 6-letter equivalents. So, in this very simple compression system you might be able to reduce the overall data storage size of your word processor documents by 30 per cent. However, compression and decompression do take some processing power so you might not use it in instances where performance is critical.



The individual benefits of flash, deduplication, and compression are enhanced if you use them together to boost performance, reduce the amount of data you have stored and the individual size of data sets to remove the need for traditional storage over provisioning strategies.

Chapter 3

Building Effective Virtualized Infrastructures Using Application-Aware Storage

In This Chapter

- ▶ Improving virtualized desktops with application-aware storage
 - ▶ Assisting with the virtualization of monolithic enterprise applications
 - ▶ Enabling more agile test-and-development environments
 - ▶ Understanding how private cloud deployments can benefit from an application-aware storage layer
-

Within the ever-growing number of increasingly virtualized information technology (IT) user cases, traditional storage technologies can and do frequently cause a variety of cost, complexity, and performance issues. This chapter takes a look at four common-use cases to illustrate how an application-aware storage layer can help to tackle and overcome such issues.

Solving Virtual Desktop Infrastructure Problems

Virtual Desktop Infrastructure, or *VDI*, is a desktop infrastructure service where each user's desktop resides within a virtual machine (VM) that lives on a server in the data center. Only the visual desktop interface is served across either a network or an internet connection, while the maintenance of the hardware, security, software, and data management is all centralized for simpler administration.



The advantages of VDI are magnified when a large number of virtual desktop users are accessing the system, such as on college campuses or in situations where the users are geographically dispersed (home workers or staff in remote branch offices with little local IT support, for example). In addition, users can access their virtual desktop and applications on multiple devices, such as laptops, desktops, tablets, and even smartphones, which helps to promote flexible working practices such as hot-desking and greater mobility.

Identifying the problems faced by traditional storage within VDI projects

VDI places a lot of demands on the *underlying storage architecture* – the architecture which maintains the data that feeds through to the desktop views served by centralized servers. This section looks at what those problems are and – most importantly – how to deal with them.

These problems can be categorized into three common concerns:

- ✓ **Lots of concurrent users:** If a class of students all decide to log onto their virtual desktops to gather material for an ongoing lecture, this process of simultaneous access can cause a *boot storm* or *login storm* which can lead to lengthy delays in accessing the VDI.

- ✓ **Unpredictable access:** Unlike other common workloads that often follow predictable process steps, each user accessing applications and data via VDI results in a random access pattern for the underlying storage. These random access patterns make tuning the storage for consistent performance difficult.
- ✓ **Over-provisioning:** To meet performance demands, VDI deployments using traditional storage technologies must allocate more resources than actually required to overcome boot storms and unpredictable access patterns. This over-provisioning is a wasteful and costly strategy that doesn't always guarantee a consistent solution.

Solving the challenges with application-aware storage



An application-aware storage layer is well suited to VDI deployments in several ways:

- ✓ One of the most important ways is **the use of flash-based storage**, instead of slower disk-based technology, which automatically serves VDI sessions to users based on an understanding of what that VDI session is doing and what other activities are taking place across the wider VDI deployment.
- ✓ Application-aware storage also includes **deduplication technology** that removes repeated and unnecessary content, along with **compression technology** that shrinks the amount of data stored even further, greatly enhancing the storage efficiency.

The combination of flash storage and data optimization techniques provides a huge performance boost that negates the issues caused by unpredictable access patterns.

- ✓ Because **the application-aware storage can look deeply into the virtualization process at a VM level (each VDI instance)**, administrators can accurately determine whether performance thresholds and service level agreements are being met. This visibility also allows the application-aware storage system to host workloads with differing performance characteristics and to dynamically allocate resources between them as required among workloads.

Checking out a real-world VDI example

South Eastern Health and Social Care Trust is part of Northern Ireland's National Health Service (NHS). The Trust provides an extensive range of quality health and social care services to over 440,000 people across Northern Ireland and employs over 12,000 staff across a range of disciplines, including nurses, midwives, social workers, dentists, and medical professionals.

A strategy to move services out into the community transformed IT requirements for the Trust, which now needed to provide remote access to all of its systems to healthcare professionals working across a wide area.

Addressing the challenge

The Trust looked at a few different technologies for providing remote access, but quickly came to the conclusion that virtual desktops were the most efficient way to deliver mobile services. The Trust had already virtualized its server environment using VMware vSphere so considered VMware Horizon View as the obvious choice for its new VDI solution. They also brought in hundreds of mobile devices, including Androids and iPads, and implemented a single sign-on (SSO) solution to ease access to those devices along with VDI.

Tackling the issue of traditional storage within the VDI project

The initial VMware Horizon View *proof of concept* (where the customer gets a unit to see the product in action) was run on the Trust's existing storage platform. 'We wanted to see how the virtual desktops would work for our remote employees,' explained Darren Henderson, Information and Communication Technology Services Manager, for South Eastern Health and Social Care Trust. 'Our existing storage devices worked OK for the average IOPS (that is, input/output operations per second) for the virtual desktops, but they were having some trouble with the spikes that come with VDI, for employee logins and other events.'

Henderson started looking for a new storage solution that was better suited to VDI whilst avoiding systems that were deemed complex to manage. 'We have a very small IT team

and don't have time to be configuring, tweaking, and fixing things all day long,' he said.

Looking at various storage solutions, and quickly realizing that an application-aware storage was what they needed, they chose Tintri – the leading application-aware storage provider.

Realizing the benefits of application-aware storage with VDI

According to Henderson, the Trust had selected Tintri because of its simplicity and performance, but also because of the relatively low cost of the systems. As Henderson explained, 'The overriding factor was the flexibility and ease of use of Tintri, VMware Horizon View, and Imprivata SSO providing our healthcare workers with access to our systems from anywhere on any device. If the environment happens to deliver cost savings as we go forward – which we believe it will – that's just an extra benefit.' He continued, 'We expect the great performance and ease of scale of Tintri storage appliances will enable us to secure the rest of our VDI project, which is expected to grow to several thousand users,' stated Henderson. 'Since Tintri is an independent node, we don't have to "size for tomorrow", we can buy just what we need for today. We don't have to buy a large, expensive frame up front, and then buy all of the disks. Tintri offers a scalable and cost-effective solution.'

The new Tintri system delivers a high performing VDI environment, with less than 13 second for initial logins and less than 3.8 second logins when switching terminals.

Virtualizing Monolithic Enterprise Applications

Some of the most critical business processes need access to complex and powerful software applications such as Enterprise Resource Management (ERP), and billing and transaction processing systems. In the past, to support these applications many organizations have deployed SAN (Storage Area Networks)-based storage technologies, using dedicated and more complex fiber-optic networks to improve performance. However, in the last five years, Ethernet-based networking

technologies have become dominant by offering both lower cost (at comparable performance levels and with simpler installation) and ongoing management. In addition, many of these *monolithic enterprise applications* (that is, single-tier software applications) are able to move away from dedicated hardware and into virtualized environments that offer the scalability and flexibility that is difficult to achieve without a continual cycle of expensive hardware upgrades.

Discovering that traditional SAN architecture causes complexity and inflexibility

Each SAN uses a system called *Logical Unit Number* (LUNs) to help configure the capacity, reliability and performance characteristics of the storage environment. A LUN is a unique identifier given to separate devices, or logical units, so they can be accessed by the networking protocol that connects the storage to the applications.

LUNs are key to the configuration of the disks (or ‘disk arrays’) used in a storage pool because these arrays are typically defined in a redundant array of independent disks (called *RAID*) to protect against failure. However, RAID is not decipherable directly by the applications, so LUNs help to manage this process and allow the assignment of pools of capacity.

Unfortunately, however, SAN and its corresponding LUN technology was created long before the arrival of virtualization. Both technologies require storage administrators to take a number of decisions – such as the number of LUNs to be created, what size they will be and what hardware the LUNs will live on – that are hard to change at a later date. This contrasts with the flexibility offered by virtualizations.

The cumbersome manual assignment and configuration of LUNs means that the underlying storage is less able to adapt dynamically to the requirements of an application, such as a surge in demand for performance or capacity. LUNs also complicate the creation of a duplicate environment to meet a disaster recovery scenario as the unpredictable nature of an IT outage doesn’t sit well with the mostly predefined nature of a SAN/LUN architecture.

Providing a virtualization layer between application and storage with application-aware storage

Rather than managing LUNs to create storage pools of a particular size and mapping multiple LUNs to changing capacity and performance needs of applications, *application-aware storage* manages a single pool of storage that offers a direct connection between the application VM, its virtual disks and the storage capacity and performance needed for them.



Application-aware storage constantly monitors the requests of each individual application and calculates and automatically assigns the capacity and performance levels from a pool of underlying storage hardware. Application-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 and not through pre-defined configuration changes.

Checking out a real-world critical enterprise application example

Makino is a global leader in the manufacture of high-performance machining centers and a provider of turnkey, engineering, and integration services. Makino manufactures machine tools that let manufacturers make specific parts for a diverse range of products. The aerospace, die/mold, automotive, medical, construction, and mining industries all use Makino's machines to produce precision components for a diverse range of products.

Addressing the challenge

Storage performance issues prevented Makino from virtualizing its business-critical applications, and managing dozens of datastores and LUNs in their existing virtual infrastructure was becoming increasingly cumbersome. The IT department looked to replace Makino's traditional SAN storage with a solution that would satisfy both the performance needs of its

business-critical applications and enable generalist administrators to manage the entire virtual infrastructure, including the storage.

Tackling the issue with traditional storage within the project

Makino had stopped short of virtualizing its business-critical applications because of concerns about storage performance. Instead, the organization relied on physical servers and direct-attached storage for running mission-critical applications, but this reliance had caused a problem. 'We had maxed out performance on our existing storage and we would have had to add much more just to meet our performance requirements for our business critical applications', said Glenn Hensley, the Makino IT infrastructure manager.

Makino was facing challenges in managing more than 40 LUNs and datastores and mapping those to hosts which forced a constant juggling of resources as its application needs changed. Makino also struggled to troubleshoot storage performance problems with its virtual infrastructure. 'Pinpointing problems involved a lot of guesswork, so any new storage solution needed to make it easier to troubleshoot VM performance and be tightly integrated with VMware', said Andy Chambers, the Makino system administrator. What Makino needed was an application-aware storage to solve the SAN and LUN management challenges in their environment.

Realizing the benefits of application-aware storage

By adopting Tintri's application-aware storage, Makino was able to successfully virtualize its business-critical applications. Makino also realized operational cost savings by increasing its server consolidation ratio. 'With one 3U Tintri VMstore system, we've been able to virtualize our business-critical servers such as SAP, SQL Server and Citrix XenApp servers', said Chambers. 'And thanks to Tintri's flash-based architecture, we get more than sufficient performance. We have moved most of our virtual infrastructure on to Tintri and still have plenty of performance available to add more VMs.'

Supporting Flexible Application Testing and Development Environments

Many organizations depend on the development of new IT process and software applications to provide them with a competitive edge or to solve challenging business problems. In order to create, modify, and test IT systems that are separated from highly sensitive production environments, software developers often need to create environments that replicate existing IT infrastructure. In many cases, developers may need to simulate a range of potential deployment scenarios, workloads, or even outages to see how new or modified applications will perform.

In the era before virtualization, creating these digital workspaces to mimic production environments required significant capital expenditure to buy identical hardware, storage, and networking systems. Significant delays would be caused while systems were bought, configured, and provisioned into the development space. For organizations with larger or even multiple development teams, the complexity of running several concurrent projects and workspaces became even more complex. The rise of server virtualization, however, has provided tremendous advantages for developers.

Thanks to virtualization, software teams can now self-provision resources on the fly and both build and test workloads efficiently by pooling infrastructure. Server virtualization has simplified the process of creating replica production environments and enabled developers to scale solutions across multiple VMs to effectively plot performance against predicted demand. Virtualization also allows the use of non-identical hardware and much better utilization of resources, including the ability to move workloads between remote data centers where computer and storage resources can be purchased on demand – a concept called *Cloud computing* (see the ‘Creating Private Cloud Deployments’ section, later in this chapter).

Understanding how traditional storage hampers modern testing and development

Unfortunately, traditional shared storage solutions don't have the same flexibility when it comes to building out virtual test-and-development environments. They tend to suffer from poor performance and management complexity in scaling to support thousands of VMs in large-scale testing and development environments. Traditional storage solutions also fail to provide consistent performance as the environments scale, which hampers developer productivity. The setup, configuration, and maintenance of traditional storage solutions also requires extensive work between the software, virtualization and storage teams, which negates many of the agile processes that modern software development techniques value.

Helping virtualized testing and development environments with application-aware storage



An application-aware storage layer is built from the ground up to support VMs and, as such, has a number of fundamental features that are particularly well suited to test-and-development environments:

- ✓ **Simple self-provisioning of storage:** Development teams can use a single interface to easily self-provision new workloads without the need to physically specify the storage placement or configuration. Administrators can provide a total pool of storage which developers can consume as required and then release when projects scale down. This process can be automated based on thresholds.
- ✓ **Reduction of storage capacity wastage from redundant data:** Test-and-development environments often have huge amounts of replicated data as teams build replica environments to test application modifications. In many cases, 99 per cent of several test environments might be identical with only minor changes to specific modules.

An application-aware storage layer helps to reduce wasted resources like these and to improve performance by using inline deduplication to discard any unnecessary physical copies of data. This feature is further enhanced by compression that shrinks the size of data as it is stored and passes across the network.

✓ **Visibility of performance and resource utilization:**

The development of new systems requires a deeper understanding of their impact on the underlying servers, storage, and network. Traditional storage systems have limited visibility into the application layer and, as such, require developers to effectively guess the underlying causes of seemingly storage-related performance bottlenecks. An application-aware storage layer gives developers much more granular visibility as to how an application – or hundreds of instances of an application – are using the storage layer. This visibility is essential for development best practice.

Checking out a real-world test-and-development environment example

F5 Networks is a global leader in application delivery that makes applications secure, fast, and available, helping organizations get the most out of their investment. F5 solutions broaden the reach of IT through an open, extensible framework and a rich partner ecosystem of leading technology and data center orchestration vendors. This approach lets customers pursue the infrastructure model that best fits their needs over time. The world's largest businesses, service providers, government entities, and consumer brands rely on F5 to stay ahead of cloud, security, and mobility trends.

Addressing the challenge

VM storage performance issues were hindering effective scaling of F5's Product Development VM environment for the company's flagship BIG-IP platform. Expanding existing storage systems to support the growing number of VMs was neither cost effective, both for initial purchase and ongoing maintenance, nor efficient. F5 wanted to deploy a VM storage solution that would not only satisfy its performance needs

but enable it to effectively scale to support thousands of VMs while simplifying management.

Tackling the issue with traditional storage within the testing and development environment

F5 was using traditional storage systems to support its test-and-development VM environment, running a combination of Windows and Linux VMs as well as its own BIG-IP Virtual Edition. ‘Performance bottlenecks in our existing storage were hindering us from deploying large numbers of VMs, affecting developer productivity. Further, we could only operate safely at less than 50 per cent performance utilization given the active-active configuration’, said Yens Jimenez Steller, manager of the F5 Product Development Lab in Seattle.

Realizing the benefits of application-aware storage

The move to an application-aware storage platform enabled cost-effective flash-based performance which supports over 1,500 VMs on just two Tintri T540 VMstore systems while allowing use of full system capacity. Data center space has also been reduced by 75 per cent compared to previous storage systems. ‘Compared to our previous storage, Tintri VMstore can accommodate twice the IOPS at less than a third of the latency in one fourth the footprint. Also, unlike our previous storage, we don’t need to spend on dedicated administrator training for managing Tintri systems.’

Creating Private Cloud Deployments

During the last decade, many organizations have transformed their IT operations in order to deliver IT infrastructure and software applications as a service, as opposed to the rigid old-school IT. This has enabled the emergence of private cloud infrastructures outside of central IT. Private cloud allows line of business owners to sidestep the slow and conservative central IT to gain the kinds of agility, responsiveness, and scale that new applications demand.

In many cases, these services are effectively commoditizing the organization's fundamental compute (or processing), storage, and networking resources and turning them into capacity charging models. A good analogy would be to compare buying and running your own oil-feed power generator to make electricity when, instead, you could buy electricity as a service from the national electricity grid.

The arrival of virtualization has led to a boom for enterprises who are able to efficiently provide their IT infrastructure and capacity in a much more granular and efficient manner to their consumers. Customers can now buy individual virtual machines complete with virtual storage and communication bandwidth purchased only based on what is required to meet the use case. This paradigm is often referred to as *cloud computing*, and it is one of the fastest-growing parts of the IT industry. Enterprises offering such paradigm within their data center are referred to as adopting a *private cloud infrastructure*.

Accepting that traditional storage isn't fit for a cloud world

Traditional storage technologies weren't designed for fast-changing environments where multiple internal customers are constantly creating and destroying workloads of different types but still sharing a pool of storage. For example, if one user needs a lot of high-performance reads for a database VM while another is running a VM with a lot of write requirements, such as data backup, traditional DAS (Direct Attached Storage), NAS (Network Attached Storage), and SAN (Storage Area Network) technologies struggle to cope with these disparate use cases.

As the requirements of each customer are often different and subject to change, the inflexible nature of SAN (with its rigid LUNs and mapping to physical servers) or the limitations around sharing data which are posed by DAS make both platforms unsuited to service providers and cloud technologies. NAS is more flexible, but still requires the creation of islands

of storage, which have capacity and performance scalability challenges. New cloud applications are sources of competitive advantage – think, for example, of the mobile apps that even the most staid enterprise companies are now using to drive engagement with their customers in a mobile world. These cloud applications are updated frequently, quickly scale up and down in response to demand, and are designed to be modular so that a single ‘application’ may actually consist of tens, hundreds, or thousands of small virtual machines that are constantly changing. Private cloud deployments require incremental performance and/or capacity upgrades to scale the infrastructure in line with the customers’ demands resulting in considerable capital and operational expenditures.

Solving private cloud deployment challenges



Unlike in a traditional-use case deployed by end-user customers, private cloud deployments need to build out products that can solve a wider range of customer use cases while meeting a number of business constraints. In this, application-aware storage can deliver (see Table 3-1).

Table 3-1 Application-aware storage versus traditional storage for private cloud deployments

<i>Private cloud requirement</i>	<i>Traditional storage issue</i>	<i>Benefit with application-aware storage</i>
Space for equipment within a data center is at a premium. We need to have as dense a footprint as possible to maximize profitability.	DAS, SAN, and NAS all have relatively low utilization rates and create islands that must be manually assigned to individual tenants within the service.	An application-aware storage layer uses deduplication and compression to remove unnecessary data and reduce the amount of wasted storage capacity. This leads to better utilization and a much denser infrastructure.

<i>Private cloud requirement</i>	<i>Traditional storage issue</i>	<i>Benefit with application-aware storage</i>
We need to be able to offer different service and performance levels for our storage layer. Doing so is essential for meeting the expectations of different customers and to enable us to sell different classes of product, like basic, standard, and premium products.	Tailoring storage to hit a predefined performance point often requires over-provisioning with traditional technologies. This method is expensive and very inefficient for a service provider with physical space constraints and tight profit margins.	Application-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. This allows the mapping of different performance bands for varying charging levels.
We need to be able to offer disaster recovery and backup services that operate on a VM level.	There are separate VM aware, data protection, and disaster recovery solutions that will work with traditional storage but they do add additional levels of complexity to management tasks.	VM-level backup and replication technologies are built directly into the application-aware storage architecture which simplifies design and allows for self-service recovery tools.
We need to accommodate many more VMs and, with them, more snapshots, clones, and policies.	Traditional storage architectures have a finite number of available LUNs and snapshots that may not scale when applications that today reside in one monolithic VM are refactored into many more small VMs.	Application-aware storage was designed to handle a much larger number of objects (VMs, vDisks, snapshots, and the like) and is able to meet cloud scale requirements far more easily.

(continued)

Table 3-1 (continued)

<i>Private cloud requirement</i>	<i>Traditional storage issue</i>	<i>Benefit with application-aware storage</i>
Private cloud environments are in a constant state of change, with VM life-times measured in hours or days instead of months or years.	Operations like storage migration of VMs to less-bust LUNs, or draconian QoS policies, do not work when the environment is in a state of constant change. It can even be difficult to determine which VM to troubleshoot when an 'application' is actually multiple VMs.	Automatic QoS policies of application aware storage give each VM its own IO 'lane'. Constant change is a given, and the storage automatically adapts to the new conditions, giving each VM the best performance automatically.

Checking out a real-world private cloud example

Northwestern University Information Technology (NUIT) is a service arm of Northwestern University that is dedicated to advancing its competitiveness, influence, and reputation. NUIT deploys, supports, and administers the IT and network infrastructure that facilitates the dynamic learning, teaching, and research activities at Northwestern University.

Addressing the challenge

Northwestern University wanted to create a private cloud infrastructure to centralize computing and storage resources. By moving to a cloud, local IT departments would be able to manage the computing and storage resources as well as the provisioning and management of virtual machines. At the time, Northwestern was using a mix of traditional fiber channel (FC) storage systems in their virtualization environment which required sophisticated SAN expertise which, in turn, forced complex direct management of LUNs and datastores.

Tackling the issue with traditional storage within the private cloud environment

To solve these issues, NUIT opted for an application-aware storage architecture. ‘We chose the Tintri T540 storage appliance as the basis for our private cloud infrastructure because of the performance it delivered in such a small form-factor while providing very simple-to-use VM granular management’, said John Walsh, manager of processing and information platform services. In tests, a single Tintri T540 appliance performed better than the FC-based arrays, ‘. . .so we don’t have bottlenecks and are able to run our most demanding database workloads’, says Walsh.

Realizing the benefits of application-aware storage

The private cloud model Northwestern deployed enables central IT staff to focus on core infrastructure management, while campus IT staff use a self-service mechanism to create VMs based on the infrastructure. ‘We streamlined the request and approval process from provisioning and usage of VM resources so everyone benefits’, said John Walsh, manager at Northwestern University’s Processing and Information Services. The Tintri approach reduced storage complexity and helped to enable Northwestern’s private cloud infrastructure deployment. ‘The single datastore per-pod model and an intuitive graphical user interface dramatically simplified administration, allowing local IT staff to monitor performance and capacity metrics on a per-VM basis for troubleshooting’, said Walsh.

Chapter 4

Understanding Key Management Concepts to Get the Most Out of Application-Aware Storage

In This Chapter

- ▶ Defining flash-centric architecture
- ▶ Creating quality of service on an application basis
- ▶ Realizing the benefits of data management at an application level
- ▶ Understanding end-to-end visibility and control

When organizations plan to benefit from application-aware storage as part of a wider move to virtualized information technology (IT), gaining an understanding of certain management concepts is critical. This chapter provides an explanation of those key concepts.

Understanding Flash-Centric Storage



Flash-centric storage offers a non-mechanical method of writing, storing, and reading digital information. Compared to a traditional hard disk drive, which uses spinning platters a bit like a compact disk to store information, flash has increased transfer performance and lower energy

requirements. Another bonus is that flash storage also offers a shorter delay in the time it takes to find the start of a section of information and transfer it to a destination. *Flash-centric storage architecture* takes advantage of these benefits of non-mechanical storage technologies to provide consistent and high performance of virtualized environments.



Flash has its downsides, however. As of 2013, the cost per gigabyte of flash storage is approximately ten times more expensive than the equivalent hard disk drive cost per gigabyte. Even though flash read performance may be an order of magnitude faster than that of a hard disk, write speeds are only between one and three times faster. Flash also suffers from a form of wear and tear that reduces performance if data is always written and read from the same location of the flash. In response, developers have come up with techniques for mitigating the perceived shorter lifespan that flash has over traditional spinning disk.

Locating flash within storage architectures

Flash can reside in different areas within storage and IT architectures to improve performance. Each location has a number of considerations, advantages, and disadvantages:

- ✔ **All-flash storage array:** In this method, all the physical hard disk drives (HDD) are replaced by flash solid state (SSD). This provides an increased performance across all-disk activity but with a huge cost premium. In some cases, an all-flash storage platform may not deliver the expected performance boost as some of the underlying technology elements supporting the storage aren't designed for the characteristics of an all-flash platform.
- ✔ **Flash storage cache:** A flash-based cache sits ahead of traditional physical storage and acts as a staging area for data reads and writes. By acting as a high-performance tier of storage that can intelligently move hot or frequently used data into the faster flash, the smaller cache can speed up perceived performance of the larger overall storage pool in a much more cost-effective fashion.

- ✓ **Server-side flash:** Placing a flash drive within the server hardware that also runs either physical or virtual applications can boost performance, providing the application is designed to utilize the local storage capacity. This method may prove to be the simplest and most effective within environments with a single digit number of servers and little need for flexibility or scale. However, this method isn't suitable for larger environments with lots of shared resources because of the increased management complexity and cost that it involves.
- ✓ **SSD and HDD hybrid virtual storage:** Although a relatively new design, some storage arrays combine flash alongside hard drives to create flash-based storage architecture for the virtualized environment. In this hybrid method, the storage appliance uses intelligence to place the data items that benefit from flash performance on flash and the less-performance dependent data on slower solid disk. This process happens constantly at the virtualization layer which boosts performance based on the demands of the virtualized application.

Figuring out how application-aware storage uses flash to provide consistent performance

Within an application-aware storage platform, flash is used effectively with deduplication and compression to provide cost-effective performance. Additional techniques, such as actively monitoring the read and write IO patterns of applications, provide consistent performance. (See the example application-aware storage from Tintri in Figure 4-1.)

The first step is *inline deduplication*, which removes unnecessarily duplicated information to ensure that the more expensive flash disks store unique items of data that applications need frequently. This data is then *compressed* to further reduce the amount of precious flash disk capacity used for storing data.

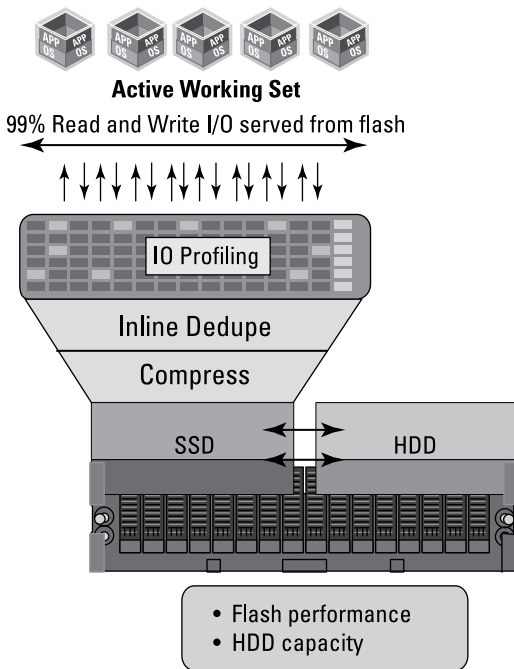


Figure 4-1: Tintri application-aware storage architecture

Creating Quality of Service on an Application Basis

Each virtual machine (VM), and the associated application that runs within it, makes demands on the underlying resources in the storage system. Some applications can be sensitive to the latency incurred by the storage subsystem while others can be sensitive to the throughput offered by the storage subsystem. As well as this, many applications have a predictable performance expectation (known as Quality of Service, or QoS) from the storage system, especially in the presence of other applications that share the same storage system.

Knowing the drawbacks of traditional storage in delivering QoS

In traditional storage, these individual actions (such as reads and writes) of each application weren't easily understood, simply because there was no visibility into the VM hosting the application and because multiple-application VMs were sharing the same LUN (that is, the Logical Unit Number) or volume. For an IT administrator to ensure a desired QoS to the application, he or she would have to manually tune the underlying storage hardware and ensure that storage associated with an application resided on a predefined area of the hardware array or else risk highly variable performance.

Highlighting key features of application-aware storage in delivering QoS



Application-aware storage is designed to automate many of the manual functions needed to define, manage, and maintain QoS at an application level. Although each storage system could provide varying management tools with different features, the most crucial include:

- ✓ Auto-configuration and performance load balancing of data across multiple application-aware storage appliances and datastores.
- ✓ Historical views of VM-level storage performance that enable organizations to meet service level agreements and create reports for more capacity and performance troubleshooting, trend analysis, and planning.
- ✓ Integration with hypervisor management tools to greatly reduce the number of management steps required to allocate and manage storage within virtualized environments.
- ✓ Management of multiple storage system appliances using a single plane of glass, to reduce complexity for IT managers.

- ✓ Seamless integration into multiple hypervisors to provide VM-level data management for data protection and disaster recovery.
- ✓ Visibility over the individual datastores and the performance or data access patterns at datastore (or groups of datastores) level.

Accruing the Benefits of Application-Based Storage Management

As well as help that application-aware storage management tools provide in delivering improved quality of service, the switch to an application-aware storage platform provides benefits to a number of data management tasks that were problematic within traditional storage architectures.

Identifying issues with data management tasks in traditional storage environments

In traditional shared storage environments, resources are allocated from the storage array based on declaring a physical group of storage resources to an LUN, which can be presented to one or more hosts for one or more uses. However, different types of workloads, applications, and data types have very different performance and protection requirements, often called a *redundant array of independent disks* (or RAID) group. Administrators would need to carve up LUNs using a complex set of rules to try and meet the protection, performance, and capacity needs across different use cases (see Figure 4-2).




Figure 4-2: Traditional storage with LUNS and different underlying RAID configurations to satisfy different application needs.

This traditional approach to management leads to several issues including:

- ✓ Inflexible allocation of performance and capacity,
- ✓ Wasteful over-provisioning to ensure peak and future performance needs, and
- ✓ A need for admin to manually reclaim space and continually tune the storage array as workloads change.

Solving traditional storage challenges with application-aware storage

Application-aware storage configurations have no LUNs or predefined RAID schemas to consider as the system uses a thin provisioning layer that abstracts the complexity of where data is physically placed on the disk (see Figure 4-3). The disk layer is also a mixture of slower spinning disk and faster flash disk, and the application-aware storage system also manages which data sits on which type of media.




Figure 4-3: No LUN complexity and common storage layout for all applications in application-aware storage.



This application-aware storage approach to management offers several other advantages:

- ✓ Capacity is automatically released as workloads are scaled back.
- ✓ Protection schemas can be based on data type or even individual VMs, and these definitions can adapt to new drivers in a non-disruptive fashion.
- ✓ Thin provisioning means that adding capacity doesn't require reconfiguration of VMs or applications, or even any downtime.

Protecting applications and data



The switch to application-aware storage offers significant benefits in terms of application protection and corresponding disaster recovery strategies. Whereas traditional storage systems require a complete copy of an entire LUN (which can contain dozens of VMs and applications) or sets of LUNs, application-aware storage systems can make rapid copies, often called *snapshots*, of just an individual VM. These snapshots offer an image of the VM at a particular point in time

and can be used to rebuild or *clone* a VM for management tasks, such as scaling up an environment with identical VM and recovering an application in the event of some form of IT outage. For data protection and recovery in the wake of any disaster or site failures, these snapshots can be sent to systems at a remote location (called *replication*), again at an individual application basis.

Because the application-aware storage understands VMs and applications, snapshotting and cloning tasks can take place directly inside the storage system without having to be orchestrated by an external agent. This setup makes data protection tasks almost instantaneous as data is not read, transmitted, and rebuilt on another platform. This provides much more flexibility in the development of protection strategies than those offered by traditional storage.

Understanding End-to-End Visibility and Knowing Why It Is Beneficial

One of the fundamental benefits offered by virtualization is the ability to simplify the creation, deployment, and ongoing management of virtualized applications through software-based tools. However, managing the critical data storage layer underpinning these processes has, in the past, required a number of discrete and separate management steps.

The common sets of tools used to manage virtualized servers are supplied by virtualization platform vendors such as VMware, Microsoft, Red Hat, and Citrix. These tools have application programming interfaces that enable the storage and networking vendors to pass information into, and to control workflows from, these virtualization management tools. This level of integration not only improves the operational efficiencies but also provides end-to-end visibility into the infrastructure and helps a lot in troubleshooting.



Traditional storage doesn't work at the same level of abstraction at which the virtualization platforms work and, therefore, fails to provide end-to-end visibility at the application level. Application-aware storage is well suited to provide end-to-end visibility as it does works at the same level of abstraction as the virtualization platform vendors.

Chapter 5

Ten Tips for a Successful Migration to Application-Aware Storage

In This Chapter

- ▶ Identifying data growth and performance
 - ▶ Building for the future
 - ▶ Calculating true cost
 - ▶ Embracing flexibility
-

If you're reading this chapter first, we're guessing it's because you're keen to avoid making mistakes that could impact a successful move to an application-aware storage infrastructure. If so, you've come to the right place.

Here are ten key issues to consider.

Set Realistic Expectations

At the outset of your migration to application-aware storage, starting with a realistic expectation of how your storage needs are likely to grow will enable you to select the most appropriate platform. So, before selecting any storage technology, conduct a basic audit of your existing platforms. Doing this will enable you to base your estimate of data growth on past trends and future strategy.

Create Benchmarks

Benchmarking the performance metrics of different application tasks (such as VDI, webserver, database request, and other common application use cases), to help define where application-aware storage can make the biggest impact, is a wise early move. After taking this step, many organizations find that they have one problem area that prompts a storage upgrade.

Grow Your Storage Capacity on Demand

Application-aware storage is a form of *scale-out storage architecture*, meaning that you don't need to buy more storage (or see more capital expenditure) than you initially need. Instead, as your capacity or performance requirements grow, you can buy additional simple-to-add appliances to scale both capacity and performance on demand.

Support Multiple Virtualization Vendors

Although VMware is the most popular virtualization software supplier, future departmental projects or even mergers and acquisitions may leave your enterprise with a mixed virtualization environment that includes Microsoft, Citrix and Red Hat products. Building in support for multiple virtualization stacks is a sensible future proofing feature for any application-aware storage layer.

Update Outdated Storage Management Policies

Moving away from complex management of LUNs (see Chapter 3) and storage arrays to an application-aware storage enables you to provision resources based on performance and capacity demands. This ability, in turn, gives you the opportunity to define better storage management processes, which can significantly reduce management costs and greatly simplify troubleshooting processes.

Include Management Time when Calculating Storage Costs

A key component to include when calculating the overall cost of your storage needs is the *operational expenditure* – the time it takes to provision, manage, and maintain the platform you use as workloads change over time. A highly automated application-aware storage system that removes the need for a full-time administrator offers considerable longer-term operational cost savings over cheaper hardware that requires lots of manually intensive tasks.

Employ IT Security Experts

Digital data is incredibly valuable. For that reason alone, a move to an application-aware infrastructure – indeed, any storage upgrade project – should include your IT security team from the earliest stages of the project.

Holistically Evaluate the Virtualized Environment Performance

Storage performance is just one of several variables that make up the overall performance of a virtualized application. The underlying data network and number of virtual servers per physical server hardware also have a massive impact on the final performance levels. Remember that claimed raw performance numbers are not the same as what a virtualized application can realize and consistent performance is very important to deliver a smooth end-user experience.

Consider Your Backup and Recovery Strategies

Application-level snapshots, clones, and replication offer much more space-efficient and cost-effective data protection and disaster recovery strategies than traditional storage arrays or tape libraries do. Migrating to application-aware storage gives you the opportunity to reconsider and enhance your strategies in these areas and to customize your policies on a per-application basis.

Seize the Flexibility Offered by Application-aware Storage

The rise of virtualization and the cloud has shown that technology is constantly evolving. A strategy that places more intelligence about the storage layer into adaptable software layers, instead of within inflexible hardware platforms, offers a major long-term advantage for IT administrators and for the overall business. Enjoy the flexibility offered by migrating to application-aware storage.



Tintri

What if your storage could see, learn and adapt?

Tintri enables IT to focus on virtualised apps instead of managing storage infrastructure. Tintri application-aware storage provides VM-level visibility, control, insight and agility, moving IT from reactive to strategic.

Discover why global enterprises such as AMD, F5 Networks, GE, NEC, NTT, MillerCoors and Time Warner choose Tintri storage.

www.tintri.com

Create a successful storage architecture for supporting virtualized apps

Datacenters are rapidly moving from a 'one application, one server' setup to a virtualized shared-server model, but storage remains the primary obstacle to accelerating virtualization growth and cloud deployments. Virtualization requires innovative application-aware storage architectures to overcome the complexity, performance, control, and cost obstacles that exist with traditional storage in virtualized IT environments and within the cloud.

Read all about it in this book – see, learn, and adapt.

- *Identify the challenges of traditional storage in a modern IT environment*
- *Understand application-aware storage*
- *Build an effective infrastructure*
- *Get the most out of your storage*

Saradhi Sreegiriraju is VP, Product Management and Technical Marketing, and has more than 16 years' experience of building enterprise storage and network management systems.

Vineet Kakani is Director of Product Management, and has more than 10 years' experience in building products in the storage and virtualization space.

Will Garside is a freelance editor, a published author, and an IT journalist with more than 17 years' experience across web, print, and broadcast media.



Open the book and find:

- Techniques for ensuring a successful virtualization project
- Best practice tips for application-aware storage
- Practical examples of application-aware storage in the real world
- What you need to know to build smart application-aware storage virtualization and private cloud

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

WILEY END USER LICENSE AGREEMENT

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