

Brought to you by **HITACHI**
Inspire the Next
©Hitachi Data Systems

Storage Virtualization FOR DUMMIES[®]

Hitachi Data Systems Edition

Maximize storage
utilization and
simplify management

**A Reference
for the
Rest of Us!**[®]

FREE eTips at dummies.com[®]



Adrian De Luca
Mandar Bhide

***Storage
Virtualization***
FOR
DUMMIES®
HITACHI DATA SYSTEMS EDITION

by Adrian De Luca and Mandar Bhide



WILEY

Wiley Publishing, Inc.

These materials are the copyright of Wiley Publishing, Inc. and any dissemination, distribution, or unauthorized use is strictly prohibited.

Storage Virtualization For Dummies, Hitachi Data Systems Edition®

Published by

Wiley Publishing, Inc.

111 River Street

Hoboken, NJ 07030-5774

Copyright © 2010 by Wiley Publishing, Inc., Indianapolis, Indiana

Published by Wiley Publishing, Inc., Indianapolis, Indiana

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, the Wiley Publishing logo, For Dummies, the Dummies Man logo, A Reference for the Rest of Us!, 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. Wiley Publishing, 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.

ISBN: 978-0-470-59770-5

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

These materials are the copyright of Wiley Publishing, Inc. and any dissemination, distribution, or unauthorized use is strictly prohibited.

About the Authors

Adrian De Luca brings over 15 years of experience in information technology to Hitachi Data Systems. In his role as Director of Storage Management & Data Protection for Asia Pacific, he is responsible for developing and managing the storage solutions business in the region, with a clear focus on storage management, virtualization, and data protection solutions. Adrian has maintained senior positions in systems engineering and software development. He holds a Bachelor of Applied Science in Computer Science from the Royal Melbourne Institute of Technology in Australia.

Mandar Bhide has over 18 years of experience in information technology. In his role with Hitachi Data Systems, as Technical Director of Storage Management & Tiered Storage Solutions for Asia Pacific, he is responsible for providing technical support for the storage solutions business and sales development, with a clear focus on storage management and virtualization solutions. Mandar has maintained positions in product management and software quality assurance in his previous roles. He holds a Bachelor of Computer Science and Masters in Management Science from Pune University in India.

Hubert Yoshida (author of the Foreword) is Vice President and Chief Technology Officer of Hitachi Data Systems. In 2006, he was named “CTO of the Year” by Jon Toigo and, in October 2006, Byte and Switch named him one of “Storage Networking’s Heaviest Hitters.” His blog was recently ranked among the “top ten most influential” within the storage industry by Network World.

Hu is a graduate of the University of California at Berkeley with a degree in Mathematics. He has authored several papers on storage area networks, fibre channel, multi-protocol SANs, and storage virtualization technologies. Hu currently sits on the Scientific Advisory Board for the Data Storage Institute of the Government of Singapore.

Authors’ Acknowledgments

Adrian De Luca and Mandar Bhide wish to thank Hitachi Data Systems’ Hubert Yoshida, David Merrill, and Ojay Bahra for their assistance with the writing of this book.

These materials are the copyright of Wiley Publishing, Inc. and any dissemination, distribution, or unauthorized use is strictly prohibited.

Publisher's Acknowledgments

We're proud of this book; please send us your comments through our Dummies online registration form located at <http://dummies.custhelp.com>. For other comments, please contact our Customer Care Department within the U.S. at 877-762-2974, outside the U.S. at 317-572-3993, or fax 317-572-4002. For details on how to create a custom For Dummies book for your business or organization, contact bizdev@wiley.com. For information about 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:

Acquisitions, Editorial, and Media Development

Senior Project Editor: Zoë Wykes

Australian Edition Editor:
Kerry Davies

Editorial Manager: Rev Mengle

**Business Development
Representative:** Karen Hattan

**Custom Publishing Project
Specialist:** Michael Sullivan

Production

Proofreader: Marguerite Thomas

Publishing and Editorial for Technology Dummies

Richard Swadley, Vice President and Executive Group Publisher

Andy Cummings, Vice President and Publisher

Mary Bednarek, Executive Director, Acquisitions

Mary C. Corder, Editorial Director

Publishing and Editorial for Consumer Dummies

Diane Graves Steele, Vice President and Publisher, Consumer Dummies

Composition Services

Debbie Stailey, Director of Composition Services

Table of Contents



***Introduction*..... 1**

What's in This Book	2
Icons Used in This Book.....	3

Chapter 1: Storage Virtualization: Getting Started 5

Defining Virtualization	5
Looking for What Characterizes an Ideal System	7
Storage Virtualization: A Stroll Down	
Memory Lane.....	7
Drivers for Storage Virtualization	10
Out of storage again!.....	11
Reducing cost	12
Managing service levels	12
Adding agility and flexibility.....	13
Repurposing and adding value to	
existing assets.....	14
Going “green”	15

Chapter 2: Types of Storage Virtualization 17

Storage Virtualization in Three Layers of	
Infrastructure	19
In the server	19
In the storage network	22
In the storage controller	26
Choosing the Right Solution	29
What are your objectives?	29
Where are you today?.....	29
Price versus cost	30
Matching growth to scalability.....	30
Can you afford an outage?	30

Chapter 3: The Economics of Storage Virtualization 31

How Do You Measure the Economic	
Benefits?.....	32
So Where Are the Savings in Storage?.....	34
Where Do You Start?	35
Hitachi's Storage Economics Model.....	36

Chapter 4: Making Storage Virtualization Real	39
Key Concepts in Storage Virtualization.....	40
Planning for the Virtualized Storage Environment	41
Implementing Storage Virtualization	42
Partitioning of storage hardware	43
Internal and external storage configurations	43
Migrating Data to a Virtualized Storage	
Environment	45
Classifying the data types	46
Creating storage tiers	46
Uninterrupted data migration	47
Chapter 5: Getting the Best from Storage	
Virtualization	49
Optimizing Performance and Improving	
Storage Utilization	49
Understanding thin, or dynamic,	
provisioning.....	50
Creating a dynamic provisioning pool	50
Benefits of dynamic provisioning	51
Management Challenges in the Virtualized	
Storage Environment.....	53
Dealing with the challenges.....	55
Managing performance and capacity	56
Chapter 6: Ten Best Practices for Deploying Storage	
Virtualization	57
#1: Choose the Right Storage for	
the Host's Mapping.....	57
#2: Assign External Storage Configurations.....	58
#3: Format the Virtual Disk	58
#4: Choose the Right Disk Attributes for	
Data Migration.....	59
#5: Migrate the Data.....	59
#6: Classify the Data.....	60
#7: Understand the Workload or I/O Profile	61
#8: Design the Dynamic Provisioning	
Storage Pool.....	62
#9: Create a Dynamic Provisioning Pool	63
#10: Plan and Design a Virtualization Project	63

Foreword



When I was a young man, fresh out of college, I could throw all my worldly possessions into the back of my Chevrolet and go wherever I needed to be. Since then, I got married, had a family, bought a house, bought furniture, appliances, and tons of other stuff to support my family's lifestyle, and now it would take several truckloads to move me with all my stuff. I have lived in the same community for the past 20 years and have built a life here. It would be difficult for me to move if my work required it. However, with the availability of the Internet and broadband, my workplace has been virtualized and I can do my work without disrupting my life and moving all my stuff.

Just like my life, businesses are burdened with tons of stuff, mainly data, that they must store, preserve, protect, and access. That burden is nearly doubling every year and is beginning to strangle businesses, limiting their ability to do their work.

Data center managers are handling aging applications that are running on aging infrastructures and facilities that were built when electricity was readily available and carbon footprints were not a concern. They spend more and more of their data center budget on maintaining old systems instead of investing in new systems that are more efficient and environmentally friendly. The reason it is so difficult to modernize IT is the burden of data that must be carried forward with the modernization of applications and infrastructures. Ten years ago, before the turn of the millennium, changes could be made by suspending operations during off hours while data was migrated or transitioned to a new storage platform. Storage frames were directly attached to application servers and capacities were small enough that they could be moved within a day.

Since then, the explosion of data, globalization, and online access requires applications to be available around the clock, and suspension of operations must be kept to an absolute minimum. The capacity of storage frames has increased a hundredfold and many more applications now share access to these large storage frames, making it difficult to coordinate downtime for shared storage frames. This makes

it nearly impossible to migrate applications or modernize the infrastructure without a major outage.

In my case, my ability to work was greatly simplified through virtualization of the workplace. I can do my work anywhere there is broadband access.

In the business world, there is a similar need to virtualize storage in order to simplify operations and make it easy to modernize IT. Only through virtualization of storage can data centers hope to have the agility to modernize IT without major disruptions to their business.

To many people, the subject of storage virtualization is very complex. The purpose of this book is to simplify this complexity, remove the misconceptions and hype around storage virtualization, and leave you with a basic understanding of what storage virtualization is and what it can do for you.

Hubert Yoshida
Vice President and Chief Technology Officer
Hitachi Data Systems

Introduction

A lot has been written, taught, and even debated about virtualization in the information technology industry over the past few years, especially in the server world where it is best known. Vendors like VMware and Microsoft actively promote how their virtualization software provides efficient resource utilization, dynamic capabilities and ease of management, while the hardware vendors fiercely contend to be the company that has the lowest power consumption, better interoperability, and ease of deployment and support.

Vendor “marchitecture” aside, all organizations are looking to improve efficiency, drive up productivity, and maximize every dollar of their IT investments. Virtualization has become an effective way to do more with less, especially in these tough economic times.

However, for all the hype generated over virtualization, there remains a considerable number of misconceptions and general confusion. Many people think virtualization is some new whiz-bang technology when, in fact, the concept has been around a long time, with its origins dating back to the mighty mainframe era. The philosophy of aggregating physical components in a coordinated way to increase systems utilization and efficiency is very mature and has been successfully applied in many areas of IT beyond servers, such as networking and applications.

So what about storage virtualization? Can the same benefits be realized when virtualization is applied to the storage area network (SAN)? The fact is that virtualization in the storage environment shares many of the same benefits of server virtualization and, most importantly, is highly complementary, extending agility and cost savings to the data assets.

As a reader of this book, you are most likely looking to understand what virtualization is in the storage world, what benefits you are likely to realize, and how best to deploy virtualization in your environment to get the most out of it. *Storage Virtualization For Dummies*, Hitachi Data Systems Edition, explains the differences in implementation among various vendors and, more specifically, how it is applied by Hitachi.

It also tries to dispel the myths to bring an objective look at how this technology is best applied.

Much like other forms of virtualization, the path is not always simple. Objectives need to be clearly understood and agreed at the outset, and a sobering understanding of where you are today needs to be discovered. In addition, as with any IT project, preparation and planning needs to be invested in the design, and the solution correctly implemented, if you are to realize the true benefits of virtualization.

The good news is that this journey does not need to be walked alone. A lot of experience has been developed over the five or so years during which this technology has made its breakthrough into storage. Many organizations have transitioned their dysfunctional legacy storage environments into highly flexible and agile infrastructures that more closely align and adapt to their business needs. Vendors have the technology, best practices, and skills to assist you, and resources such as this book help put you on the right track to succeed.

What's in This Book

This book comprises six chapters.

Chapter 1 sets the framework of storage virtualization, giving you a summary of what it is and how you apply it, as well as an outline of how it came about. This chapter also analyzes the drivers of storage virtualization (why you'll even think of applying it).

In **Chapter 2** we talk about the types of storage virtualization you can choose from and how to select the right type for your situation. **Chapter 3** looks at the economics of storage, including an outline of Hitachi's Storage Economics model and a case study.

Chapter 4 delves into more detail on planning and implementing storage virtualization according to your own objectives. This chapter also examines how to migrate your data to the virtualized storage environment. And in **Chapter 5** we explain how to optimize and manage the system's performance.

Last but not least, no *For Dummies* book would be complete without a chapter of tens. So **Chapter 6** quite succinctly explains the ten best practices for physical configurations, data migrations, and thin, or dynamic, provisioning, and tips for designing your first storage virtualization project.

Icons Used in This Book

Throughout this book we use a series of icons in the margins that help flag special information. Here's what to look for.



When we tell you something about storage virtualization that bears remembering, we mark it with the Remember icon. Descriptions that appear beside this icon are worth storing away because they help build your understanding.



This flags a shortcut or easy way to do something.



This icon indicates the information is really for the propeller spinners; you don't need to know all about how this works but you may find it interesting.



This icon indicates a resource on the Internet you can look at to get further information.

Chapter 1

Storage Virtualization: Getting Started

.....

In This Chapter

- ✓ Understanding virtualization and how it applies to storage
 - ✓ Identifying what makes an ideal system
 - ✓ Familiarizing yourself with the evolution of storage virtualization
 - ✓ Finding out what drives the need for storage virtualization
-

Virtualization is one of the biggest buzzwords of the technology industry right now. But what is virtualization? More importantly, what can virtualization do for you and your company?

In this chapter we look at just what virtualization is and how the concept applies in a practical sense to computer storage systems. We also take you on a stroll down memory lane to understand how storage virtualization came about and some of the business considerations that create the need for storage virtualization.

Defining Virtualization

If you look on the Internet, you'll find any number of definitions for virtualization as it relates to IT.

A general definition of virtualization is:

The act of integrating one or more (back-end) services or functions with additional (front-end) functionality for the purpose of providing useful abstractions.



Abstracting the physical components of computing resources optimizes the way in which other systems, applications, or end users interact with them. This definition summarizes two very important characteristics of virtualization: It should reduce the complexity when compared with managing devices discretely, as well as add greater capability to improve services — a bit like a one-plus-one-equals-three equation.

When described in terms of storage, the Storage Networking Industry Association (SNIA), the independent body of the storage industry, has a more specific definition (two of them, in fact) for storage virtualization, from its technical tutorial on virtualization:

1. *The act of abstracting, hiding, or isolating the internal functions of a storage (sub)system or service from applications, computer servers, or general network resources for the purposes of enabling application and network independent management of storage or data.*
2. *The application of virtualization to storage services or devices for the purpose of aggregating, hiding complexity, or adding new capabilities to lower-level storage resources.*

Put simply, storage virtualization aggregates storage components, such as disks, controllers, and storage networks, in a coordinated way to share them more efficiently among the applications it serves. Pooling these resources in a logical way helps remove the physical barriers that would otherwise exist and maximize the full potential of these resources. Hiding the complexity associated with all the components required to deliver storage helps simplify the operations of managing the environment and, more importantly, provides greater flexibility to meet the needs of applications.

Think of virtualization in the context of using a telephone. When you pick up the handset or initiate a call on your mobile phone, as the caller you are completely unaware of how you are connected to the other party. The complexity of establishing communication is completely invisible; you are oblivious as to the number of hops across the telephone network or the type of communication involved (such as satellite or underground cable), or even if there is a failure in the network. This is virtualization technology at work, using available resources as efficiently and intelligently as possible to deliver a service.

Looking for What Characterizes an Ideal System

Many vendors offer storage virtualization solutions, each implemented in a different way or offered in different packaging. Therefore, it is important to understand the characteristics of an ideal storage virtualization solution.



A good storage virtualization solution should:

- ✓ Enhance the storage resources it is virtualizing through the aggregation of services to increase the return of existing assets.
- ✓ Not add another level of complexity in configuration and management.
- ✓ Improve performance rather than act as a bottleneck in order for it to be scalable. *Scalability* is the capability of a system to maintain performance linearly as new resources (typically hardware) are added.
- ✓ Provide secure multi-tenancy so that users and data can share virtual resources without exposure to other users' bad behavior or mistakes.
- ✓ Not be proprietary, but virtualize other vendor storage in the same way as its own storage to make the management seamless.

Storage Virtualization: A Stroll Down Memory Lane

With the introduction of commercial computers back in the 1950s, like IBM's mainframes, the magnetic disk storage used to maintain information was directly attached via a cable, or a *bus*, to the computer's central processing unit (CPU) and volatile random access memory (RAM). This method of attachment, commonly referred to as *direct attached storage (DAS)*, provided a modest amount of non-volatile storage to a single computing system.

Over the course of the next few decades, standard *protocols* (the language used to communicate between devices) emerged, such as the small computer system interface (SCSI), making it easy to

connect different vendor storage to the systems. This method of connectivity extended beyond traditional magnetic storage to devices such as CD-ROMs, tape drives and autoloaders, and JBOD (*just a bunch of disks*). Although the different types of storage flourished and fault-tolerant designs provided greater reliability, their connectivity was still confined to single servers or workstations, limiting the utilization of the media.

The first wave of storage virtualization appeared in the cache controller arrays. Large numbers of magnetic disks were pooled together and arranged in a way to provide fault tolerance to protect against individual disk failures. This system became known as *RAID (redundant array of independent disks)*. Utilizing a common pool of cache memory, applications could work with a logical image of a data block rather than working with the actual data block as it spun around on a disk platter. This improved performance by masking the seek and rotation delays of a mechanical disk. It also enabled mainframes to use lower cost disks.

In the early 1980s, in the midst of the computer network revolution that allowed servers to talk to one another and exchange information over standard mediums like ethernet, vendors such as Novell and Sun Microsystems figured that, if servers can share data, then why not share their storage as well? By harnessing the power of the network, *network attached storage (or NAS as it is commonly known today)*, allowed servers to treat storage in another room, or potentially another city, as though it was right there, physically attached to the server. This offered unprecedented levels of flexibility to store data centrally but make it available widely. It also allowed multiple users to write and read to the storage at the same time, providing universal access to the same data sets and increasing collaboration between users. As this common way to access data became popular, other operating system vendors started offering this capability in their systems, and companies like Network Appliance, EMC, and Hitachi Data Systems packaged the functionality into appliances for easy deployment and increased scalability.

The next advancement in storage virtualization came with the introduction of *storage area networks (SANs)* in the mid to late 1990s, when the concept of storage resource aggregation made it out of the box.



What is a storage area network?

A *storage area network (SAN)* is an architecture to connect detached computer storage devices, such as disk arrays, tape libraries, and optical jukeboxes, to servers in a way that the devices appear as local resources. SANs deliver storage to servers at a block level, and feature mapping and security capabilities to ensure only one server can access the allocated storage at any particular time.

The protocol, or language, used to communicate between storage devices and servers is *SCSI* (small computer system interface).

Traditionally, SANs have used optical connectivity called *fibre channel (FC)* due to its gigabit transfer speeds and ability to work over long distances. However, more recently ethernet-based storage networks have become popular due to the ubiquity (and therefore lower cost) of the technology in the TCP/IP world (transmission control protocol/Internet protocol) and the adoption of the SCSI protocol, called *iSCSI* (Internet small computer system interface) to the medium.

SANs helped organizations consolidate their storage assets to improve capacity utilization by sharing their storage resource effectively. This simplified management by using common software tools, and enabled replication of critical information over long distances to provide greater levels of protection against data corruption and disaster events. Many larger organizations, typically financial institutions and telecommunication providers, were among the first to implement SANs.

Even though these improvements helped organizations extract better value from their storage investments, a number of problems still remained. Islands of SANs were created due to poor interoperability between different vendors and devices, utilization levels were still relatively low due to traditional storage allocation practices being used, and the lack of mobility or replication between different storage vendors meant organizations were still hamstrung.



What is network attached storage?

Network attached storage (NAS) has traditionally been a single device that contains disk storage and computer components (CPU, memory, network ports) and whose sole purpose is to deliver storage to servers. NAS is sometimes referred to as a ‘file server’ due to its ability to share a common pool of files with multiple servers. By managing the disks, file systems, and volumes, it optimizes and simplifies the delivery of storage to servers in a network environment, and can provide services such as redundancy and replication, relieving servers of these tasks.

NAS storage typically connects to an organization’s TCP/IP network via ethernet (the same used by servers) and uses specific languages, or protocols, like NFS (network file system) or CIFS (common Internet file system), based on the server message block protocol, to exchange information between servers.

More recent implementations of NAS appliances use SANs to access attached block-based storage.

In response to these challenges, in the early years of the 2000s, storage vendors began to introduce advanced virtualization features into their products. These features went beyond utilization improvements, also providing external connectivity of heterogeneous storage, non-disruptive data migration and mobility features, enhanced business continuity, logical partitioning, tiered storage, and thin provisioning. With *thin provisioning* (explained further in Chapter 5), the physical capacity of a volume or file system is allocated when applications write data, rather than being pre-allocated at the time of provisioning. Although the need for these capabilities was well understood by vendors, each vendor had a different philosophy on where the virtualization should reside.

Drivers for Storage Virtualization

There is no denying that advancements in computer technology have redefined methods of communication, creativity, and business management. This revolution has created an explosion of information generated for both businesses and home users. In the business world, applications that manage emails and Web

sites, as well as customer relationship management (CRM) and sales systems, are all critical to keeping the doors open. If these applications are unavailable for whatever reason, businesses can't service their customers and they'll go elsewhere. The Web has emerged as a source of rich media (images, video, and audio) and continues to grow unabated, with services like Facebook, Flickr, and YouTube creating unprecedented communities and collaboration between people all over the world. This exponential growth, together with the need to access information, has placed huge demands on the IT staff looking after them, especially on their storage infrastructures.



With this explosion of new applications and services, the value and importance of the information generated has become more vital. Business intelligence and decision support systems mine information in databases to create valuable knowledge that helps organizations make prudent decisions and stay competitive. Also, with regulatory compliance coming into play, the need to retain and protect corporate data for predetermined periods of time is now inscribed in law. Similarly, the accessibility of information and media from the Internet and Web sites, in particular its availability, determines how successful a business is. If users can't get what they want, when they want it, from their browser in a couple of seconds, then they simply go elsewhere.

All this creates several challenges in the storage environment.

Out of storage again!

Perhaps the most nerve-wracking part of a storage administrator's job is when he has to go "cap in hand" to the IT manager or CIO and ask for money to buy more storage. Like needing to ask your parents for more pocket money, and having to first explain where your weekly allowance went, system administrators face a similar judgment. Unfortunately, explaining "who took all the storage" is not necessarily simple, it just seems to vanish!

The main reason this happens is applications are typically allocated more storage than they need. Application and system administrators often request as much storage as they can get, rather than what they're likely to need, since it is a scarce resource. Twelve months down the track they find they have used only a small portion of their allocation, say 30 percent. Due to traditional storage provisioning methods, the unused

70 percent remains “stranded.” Although it hasn’t been used by the application and is free, it can’t be allocated to other applications as it is dedicated to the original server. Multiply this situation out to tens or hundreds of servers and applications and it’s not hard to see how this inefficiency quickly adds up.

Compounding the problem is the common practice of creating multiple copies of data. For every byte of production data, organizations create any number of replicas for backup and recovery, disaster recovery, data mining, development and testing. Many organizations exposed to industry or government regulations retain several more copies for compliance or other records management requirements long after the data has become stale.

Countless surveys and assessments of storage capacity utilization have been conducted over the years, most revealing that organizations suffer from very poor actual utilization, as low as 30 to 40 percent.

Reducing cost

As organizations deploy more and more applications into their environment, with consequent growth of data (often exponentially), many find they need to buy more storage more frequently, putting a real strain on budgets. As more storage is added, without effective storage management tools and processes, more IT people are needed to keep it going, and an unending cycle perpetuates. As more allocation requests are met, more people are needed to monitor the health of the infrastructure and protect the critical information by ensuring it is replicated and can be recovered in the event of a failure, requiring more storage.



With IT budgets and headcount staying flat or in many situations actually decreasing, organizations need to find more efficient and automated ways of managing their valuable information assets in the face of static or shrinking resources.

Managing service levels

IT systems have become the lifeblood of most organizations in this age. With the dependency on these applications to service day-to-day business activities, making sure they deliver consistent and reliable service is critical.

Take email, for example. Most employees take this service for granted as a fast way to communicate between their departments and with customers, send out information or quotations, and approve business decisions. If the email server were to slow down or be unavailable for a couple of days, not only would employees be unhappy, but the organization's productivity would plunge to a standstill. In this highly connected, real-time world, having to resort to traditional forms of communication and paper-based processes would significantly reduce the number of transactions the organization could do and, if the unavailability persisted, could place the business in real jeopardy. Therefore it is imperative that IT staff monitor the health of their infrastructure end to end, including storage, as the key piece to delivering dependable service levels to the business.

Adding agility and flexibility

With such demanding service levels that require applications to be available 99.99 percent of the time (or more!), bringing an application down for servicing and maintenance, or migrating to new server or storage platforms, would be unthinkable. Try telling the salespeople in your office that you have to shut down their quoting and order-entry system for a day and they're likely to go after you with pitchforks!

Therefore any such maintenance activities must be done with the least or, in many cases, no disruption. This poses a real challenge for IT staff as it typically involves coordinating multiple people to orchestrate the changes, having sound processes to execute the tasks with no data loss, and having fall-back measures in case something goes horribly wrong. All this sounds like a death-defying high-wire circus act reserved only for the brave. Many organizations try to avoid such activities as it places them in a risky situation, but migrations, for example, are simply unavoidable. Organizations need to look for infrastructure that is more flexible and agile to cope with these requirements.

In the server world, running out of CPU or memory resources to meet a given workload is common. With server virtualization software like VMware, applications and their operating systems can be seamlessly moved to other physical servers without the application skipping a beat.



For the storage infrastructure, if the storage serving an application can't provide adequate performance or availability, traditionally system and storage administrators would stop the application, copy all the data to another location, verify it and hope that the application comes back in good shape. Storage virtualization technology provides the ability to move data from one pool of disks to another without any need to do a manual migration.

Repurposing and adding value to existing assets

It's funny how things just don't seem to last as long as they used to. Not so much because of wear and tear but because they become superseded by new models with better performance or features.

Most storage subsystems and associated networks cost thousands or even millions of dollars. So, with such potentially short life cycles, it's not surprising that organizations are looking to squeeze every last drop out of their investments.

For many organizations, it's not a matter of technology being superseded, but no longer meeting their changing requirements. Most IT departments assess their storage requirements according to the most appropriate features, technology, vendor, and budget to serve their purposes best today. However, business landscapes are constantly transforming, and unforeseeable events can change things in a heartbeat. Certain applications can become more important, requiring greater protection or performance. Your organization could acquire another company, and you inherit another IT infrastructure to manage and merge with your own.

The problem is that business and therefore technology requirements change rapidly, often within a year, while hardware may be capitalized in three to five years. If you've just bought a new storage unit and are only one year into a five-year capitalization cycle or five-year lease, you're locked out of using new storage-based technologies for the rest of the capitalization or lease period.



When organizations compare the cost of maintenance for their aging hardware with that of buying new hardware, they are often motivated to buy new. However, with storage virtualization solutions you can breathe new life into existing assets, rather

than replacing them, and repurpose them for a job with fewer requirements, or extend functionality, such as thin provisioning or remote replication (these concepts are explained in Chapter 5), to increase their value to the business.

Going “green”

More than 20 years of research into carbon pollution now strongly suggests that human activity is having a profound impact on the Earth’s climate. The almost indisputable evidence has now forced many governments to action by implementing tangible strategies and policies that reduce carbon emissions. As one of the many measures to tackle the problem, it calls for countries to implement carbon emissions trading schemes to steadily reduce pollution, but at a cost to businesses.

So what has all this got to do with storage? The net effect of these new trading systems, coupled with the fact that the world’s reserves of non-renewable resources are rapidly depleting, means the cost of power is rising. In the past, the cost of powering and cooling data centers was negligible and was often not even factored into the IT budget. However, according to an IDC article in 2008, recent surveys suggest that with continued deployment of applications and data doubling every 18 months, power and cooling costs will grow eightfold by 2010! In another survey, by Gartner Consulting, findings showed that the power consumption required to run and cool data centers accounts for almost a quarter of global carbon dioxide emissions from the information and communication technology sector. And the StorageIO Group reports that storage itself accounts for between 37 and 40 percent of total energy usage from hardware.

With business accounting for a substantial share of greenhouse gas pollution, many organizations have developed social responsibility policies to help protect the environment. Demonstrating to customers and their respective industries that they are not only conscious of the issue, but care about being part of the solution, is important for organizations. A number of large multinational companies, such as Citigroup, Vodafone, HSBC, and British Telecom, have already committed to carbon reductions within their companies, some aiming to be carbon neutral altogether.

All these factors are prompting businesses to look at innovative ways to become more efficient users of power in order to comply with government policies, control their budgets, and protect their reputations, along with the environment.

Many IT vendors are already actively delivering products that have less environmental impact. Through streamlining production processes, many vendors have eliminated waste, requiring less carbon-dependent power. Smarter designs of components such as processors and power supplies, and the ability to spin down inactive disks, reduce the running costs, and use of non-hazardous materials in the product itself protects the environment when it is disposed of.

However, these initiatives only go part of the way.



With data storage growing at such exponential rates, together with the inefficiency of traditional storage networking, provisioning and application consumption methods, storage virtualization can go the next step by eliminating unnecessary waste.



Consolidating storage through intelligent pooling of storage resources and use of thin provisioning can delay storage purchases, resulting in far fewer unused disks spinning and consuming power. By tiering storage resources and aligning the right storage characteristics to the value of the information, organizations can use platforms that utilize disk spin-down technologies for archive data, for example, that is infrequently accessed.

Chapter 2

Types of Storage Virtualization

In This Chapter

- ✓ Examining storage virtualization throughout the infrastructure
 - ✓ Knowing how to select appropriate storage solutions
-

In today's modern computer architectures, virtualization exists in almost every layer, from the application to the operating system, server, networks, and storage devices. For example, application clustering technologies such as Microsoft Clusters and Oracle's Real Application Clusters (RAC) manage the process of selecting a server to deliver an application without the user knowing which server it's coming from. Server virtualization allows you to run multiple operating systems on the same physical hardware platform to improve utilization of the central processing unit (CPU) and memory. Most of these forms of virtualization work together to optimize efficiency throughout the layers of technology.

This chapter examines the different forms of storage virtualization and helps you to choose the appropriate type of storage virtualization for your environment.



Virtualization can also be implemented in various layers of the storage infrastructure, starting at the operating systems device file through to the host bus adapter, storage network, and storage array, as shown in Figure 2-1.

Virtualization concepts can also be implemented within the devices themselves. For example, storage array vendors employ virtualization aspects within devices such as RAID (redundant array of independent disks). RAID groups together physical

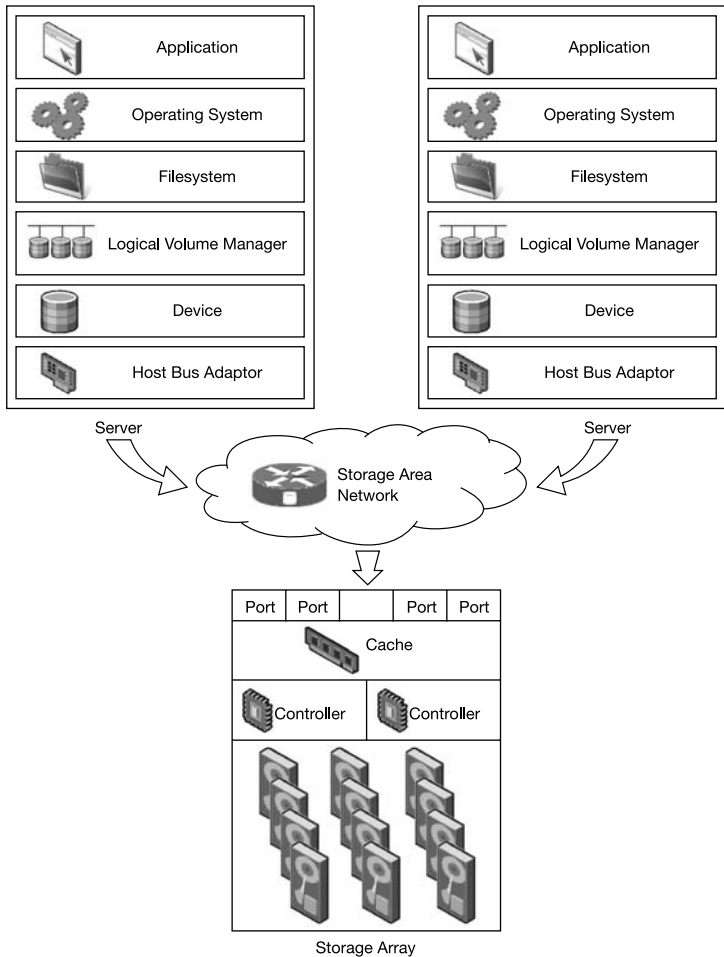


Figure 2-1: Storage architecture follows a logical system.

hard disks, spreading the data across all the spindles and slicing them up in chunks called LUNs (logical unit numbers) to deliver to servers. This method provides greater protection against a physical drive failure, so if one were to stop working for whatever reason, the system would not lose any valuable data and continue to work. Similarly, storage network vendors incorporate partitioning capabilities such as zones and virtual storage area networks (VSANs) to separate workloads for management and security purposes.

Storage Virtualization in Three Layers of Infrastructure

Modern storage virtualization technologies pool heterogeneous storage vendor products together in a specific way to provide advanced features such as non-disruptive migration of data and thin provisioning (explained in Chapter 5). This level of abstraction can be implemented in three layers of the infrastructure, in the server, in the storage network, and in the storage controller.

In the server

Some of the earliest forms of storage virtualization came not from the storage infrastructure, but from within the server, or, more specifically, the server's operating systems.

With traditional storage hardware devices that connected directly to servers, the actual magnetic disk was presented to servers and their operating systems as LUNs, where the disk was arranged into sectors comprised of a number of fixed-size blocks. To allow applications to not only store, but find information easily, the operating system arranged these blocks into a "file system." Much like a paper-based filing system, a file system is simply a logical way of referencing these blocks into a series of unique files, each with a meaningful name and type so they can be easily accessed. For example, take the file name `My_Summer_Break.doc`. `My_Summer_Break` describes what is contained within the file and the extension `.doc` identifies the file as a document. This naming is a lot more meaningful than just a numeric block number.

Although file systems helped to reference information easily, as more and more of them were created, exhausting the storage space of the physical LUN, another LUN would be created and given to the operating system to continue storing files. To know which data was stored on what LUN, the operating system would assign each one a volume number, name, or identifier. In Microsoft Windows you are most likely familiar with a letter given to each volume such as `C:\` or `D:\`, whereas in UNIX these look like `/dev/hd0` or `/dev/hd1`. As applications and users created more files, more volumes were needed to keep up, pretty soon making it very difficult to manage.

Then operating system vendors came up with the concept of a *logical volume manager (LVM)*. Much like how file systems grouped blocks together to present files, LVMs grouped volumes or LUNs together to present larger, more flexible storage pools to applications, as shown in Figure 2-2. When an LVM started to run out of space, you could *concatenate*, or add another volume, to make it larger without having to reconfigure the application or shut it down. Conversely, if you had a large volume that you wanted to slice into smaller chunks, LVMs would let you partition it in order to separate different information, say one slice for the operating system itself and the other for user data.





Figure 2-2: Logical volume managers (LVMs) store information in file systems.

Also, since LVMs could hide or abstract where the storage was actually coming from, you could present multiple physical disks and perform a “striping” operation. The technique of *data striping* takes multiple physical disks, ideally with their own discrete path to the server, breaks up the files into even pieces, and spreads the file across multiple storage devices. Writing and reading data in this way reduces the response time and thereby

increases the performance. All hard disks have a *seek time*, which is the time it takes for the information being requested on the spinning disk to arrive under the reading head. By spreading data across multiple disks, the process of finding the information can be done in parallel, thereby reducing the seek time.

Although striping greatly helped to retrieve data faster, it did pose a significant risk that, if just one of the disks failed, then effectively the full data could not be retrieved, resulting in corruption. In the 1980s, a new striping method was introduced, RAID, which increased data reliability if disks failed. RAID-1 duplicates all the data across every disk, providing a backup copy for everything. Although this provides high levels of resilience, it is extremely expensive because you need double the amount of storage. RAID-5, shown in Figure 2-3, combines the performance of striping plus parity, which provides recovery in the event of a disk failure. By using a mathematical formula to calculate the parity, if a disk were to fail in the RAID set, then the data can be regenerated and retrieved with no impact. More recently, RAID-6 was introduced to cope with the long regeneration times of large disk drives, writing parity information to two disks instead of one.



RAID-5 Set

Figure 2-3: How RAID increases resilience to minimize data loss.



Server-based forms of storage virtualization were originally incorporated into operating systems as software and still remain very popular today. Here are the key benefits of this approach:

- ✓ Server-based storage virtualization is highly configurable and flexible since it's implemented in the system software.
- ✓ Because most operating systems incorporate this functionality into their system software, it is very cheap.
- ✓ It does not require additional hardware in the storage infrastructure, and works with any devices that can be seen by the operating system.

A number of downsides to server-based virtualization also exist:

- ✔ Although it helps maximize the efficiency and resilience of storage resources, it's optimized on a per-server basis only.
- ✔ The task of mirroring, striping, and calculating parity requires additional processing, taking valuable CPU and memory resources away from the application.
- ✔ Since every operating system implements file systems and volume management in different ways, organizations with multiple IT vendors need to maintain different skill sets and processes, with higher costs.
- ✔ When it comes to the migration or replication of data (either locally or remotely) it becomes difficult to keep track of data protection across the entire environment.

Most operating system vendors like Microsoft, IBM, Hewlett Packard, and RedHat (Linux) provide at least some capability to virtualize storage resources. Vendors like Symantec offer more advanced forms of server-based storage virtualization.

In the storage network

With the introduction of network attached storage (NAS) and storage area networks (SANs) in the late 1990s, it became possible to separate disks (and their controllers) from servers and share the storage resources more effectively among all applications in the IT environment. The storage network became the traffic cop for all information being exchanged between servers and storage devices, and some storage vendors thought this would be the perfect place to manage virtualization.



Network-based storage virtualization embeds the intelligence of managing the storage resources in the network layer, abstracting the view of real storage resources between the server and the storage array, either in-band or out-of-band.

The *in-band* approach, sometimes referred to as *symmetric*, embeds the virtualization functionality in the I/O (input/output) path between the server and storage array, shown in Figure 2-4, and can be implemented in the SAN switches themselves or in specialized appliances. All I/O requests, along with the data, pass through the device, with the server interacting with the

virtualization device, never directly with the storage device. The virtualization device analyzes the request, consults its mapping tables, and, in turn, performs I/O to the storage device. These devices not only translate storage requests but are also able to cache data with their on-board memory, provide metrics on data usage, manage replication services, orchestrate data migration, and implement thin provisioning.




Figure 2-4: In-band network storage virtualization is embedded in the I/O path.

The *out-of-band* approach, sometimes referred to as *asymmetric*, does not strictly reside in the I/O path like the in-band approach; rather, it works hand in hand with specific virtualization-enabled SAN switches to perform specific look-ups, shown in Figure 2-5. The servers maintain direct interaction with the storage array through the intelligent switch. The out-of-band appliance maintains a map (often referred to as *meta-data*) of all the storage resources connected in the SAN and instructs the server where to find it. In this two-step process, the server uses

special software or an agent, as instructions need to be sent through the SAN to make it work. As data never passes through the virtualization device, performance is only slightly impacted; however, functions such as caching of data are not possible.



Both in-band and out-of-band approaches provide storage virtualization with the ability to:

- ✔ Pool heterogeneous vendor storage products in a seamless accessible pool.
- ✔ Perform replication between non-like devices.
- ✔ Provide a single management interface.

However, only the in-band approach can cache data for increased performance.

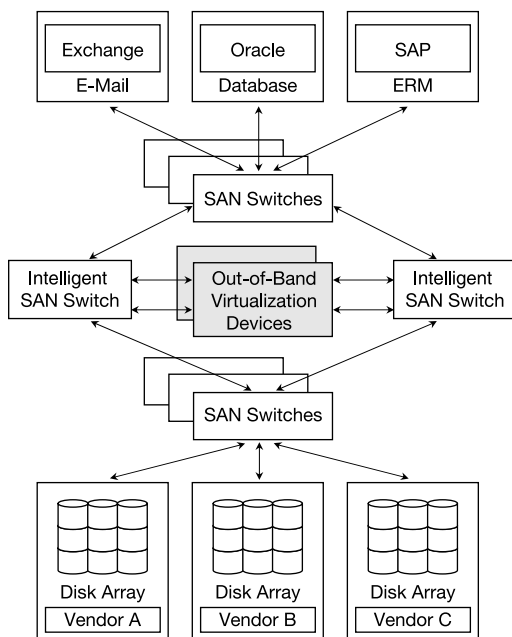


Figure 2-5: Out-of-band network storage virtualization uses intelligent switches to maintain direct interaction between the servers and the storage arrays.

Both approaches also suffer from a number of drawbacks:

- ✔ Implementation can be very complex because the pooling of storage requires the storage extents to be remapped into virtual extents. This requires a mapping table, which becomes a single point of failure, and a vendor lock-in. Migration to a different virtualization solution is extremely difficult or impossible as all the data must be moved back through the mapping table.
- ✔ The virtualization devices are typically servers running system software and requiring as much maintenance as a regular server. Clustering is needed to protect the mapping tables and maintain cache consistency between the nodes, which can be risky. Servers are also limited in the number of storage ports they can support and RAM (random access memory) capacity (as storage is limited in cache). As a result, servers lack scalability.
- ✔ The I/O can suffer from latency, impacting performance and scalability due to the multiple steps required to complete the request, and limited to the amount of memory and CPU available in the appliance nodes.
- ✔ Decoupling the virtualization from the storage once it has been implemented is impossible because all the meta-data resides in the appliance, thereby making it proprietary.
- ✔ Solutions on the market only exist for fibre channel (FC) based SANs. These devices are not suitable for Internet protocol (IP) based SANs, which utilize iSCSI (Internet small computer system interface), NAS, or mainframe servers.
- ✔ Since both approaches are dependent on the SAN, they require additional switch ports, which involves additional zoning complexity.
- ✔ When migrating data between storage systems, the virtualization appliance must read and write the data through the SAN, check status coming back, and maintain a log for any changes during the move that impact performance.
- ✔ All of these add complexity and cost to the SAN, which makes it very difficult to manage the network with *registered state change notifications* (RSCNs, which notify specified nodes of any major fabric changes), inter-switch chatter, zoning changes, and buffer credit management. SANs are the third leading cause of application failure after human and software errors.

The in-band approach has a couple of additional disadvantages:

- ✔ Specialized software needs to be installed on all servers, making it difficult to maintain.
- ✔ It also requires specialized SAN switches that support the propriety virtualization protocol, forcing organizations to replace existing SAN fabric, which can be expensive.

For all these reasons and more, storage network virtualization solutions have had limited traction in the marketplace. Vendors such as IBM, EMC, and FalconStor provide network-based storage virtualization solutions.

In the storage controller

Enterprise-class storage arrays, which have features and capability suitable for large organizations, have always featured virtualization capabilities (some more than others) to enhance the physical storage resource. One example of this is RAID, for providing data protection from disk failures. Host storage domains are another example, for virtualizing front-end connection ports to allow multiple operating systems to use the same physical port. Many enterprise-class devices incorporate sophisticated switching architectures with multiple physical connections to disk drives to provide balanced performance and resilience, as well as the ability to logically segment internal global cache for service quality and security.

With various forms of virtualization already incorporated into storage controllers for the past 15 years, some storage array vendors thought it was logical to extend these capabilities to storage resources outside the box. Controller-based storage virtualization was born (see Figure 2-6).

This breed of storage array allows other heterogeneous vendor storage arrays to be directly connected to its controllers. The external storage assets presented to it are then “discovered” and managed in the same way as internal disks. This approach has a number of benefits, including not requiring a remapping of LUNs or extents, meaning no additional layer of management and far less complexity in the network. Once virtualized in this manner, the sophisticated microcode software that resides on the storage controller presents the external storage assets as if they resided within the array, with hosts none the wiser of where they are physically connected.

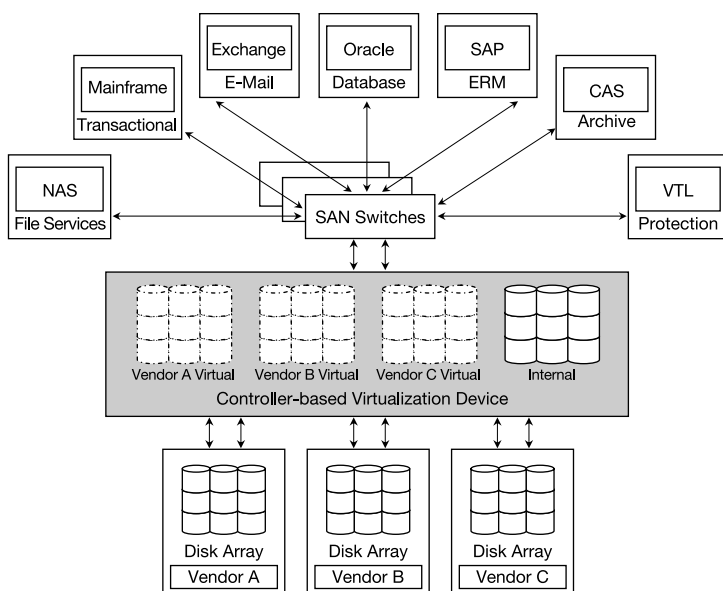


Figure 2-6: Controller-based storage virtualization allows external storage to appear as if it's internal.

This block-based virtualization effectively enables all capabilities of the virtualized storage controller to be extended to the external storage as well. Leveraging mature enterprise-class features, data can be migrated non-disruptively from one pool to another, and replication can take place between non-like and like storage. In addition, partitioning can be implemented to allocate resources such as ports, cache, and disk pools to particular workloads to maintain quality of service and security, and thin provisioning can reduce storage allocation to other assets. More recently, vendors have introduced clustering capabilities between controllers to provide the ultimate high availability in the unlikely event of a complete storage array failure.

Although a few downsides to controller-based virtualization exist, the advantages not only far outweigh them but they also address most of the deficiencies found in server- and network-based approaches.



Here are some of those advantages:

- ✔ Connectivity to external storage assets is done via industry standard protocols, with no proprietary lock-in.
- ✔ Complexity is reduced as it needs no additional hardware to extend the benefits of virtualization. In many cases the requirement for SAN hardware is greatly reduced.
- ✔ Controller-based virtualization is typically cheaper than other approaches due to the ability to leverage existing SAN infrastructure, and the opportunity to consolidate management, replication, and availability tools.
- ✔ Capabilities such as replication, partitioning, migration, and thin provisioning are extended to legacy storage arrays.
- ✔ Heterogeneous data replication between non-like vendors or different storage classes reduces data protection costs.
- ✔ Interoperability issues are reduced as the virtualized controller mimics a server connection to external storage.
- ✔ All virtualization capabilities can be extended to other forms of storage and platforms such as NAS, iSCSI, content archive, and virtual tape library through the connection of appliance heads. Mainframes can also be connected, reducing islands of stranded storage. This enables one common management, data protection, and search methodology across the entire environment.
- ✔ Interconnecting legacy assets and non-disruptive data migration from one platform to another without any application downtime assists in technology refresh cycles.
- ✔ Infrastructure can be de-virtualized with full data integrity since no meta-data is maintained in appliances.
- ✔ With virtualization embedded in the controller, an enterprise storage array has ten times the cache size of a SAN controller and more port connectivity than a SAN-based virtualization appliance. This enables it to scale well beyond SAN-based appliances.
- ✔ A storage controller can enhance the performance of attached modular storage by as much as 30 percent. SAN-based virtualization can't enhance the performance of external storage and in most cases will degrade it.

Vendors such as Hitachi Data Systems, Sun Microsystems, and Hewlett Packard employ this form of storage virtualization.

Choosing the Right Solution

Storage virtualization is an “enabler” in your IT infrastructure; it helps you maximize the potential of storage assets, creating efficiencies and automation, and helps avert risky situations. But how do you figure out which approach is the right solution for you?



Consider the following issues and how they relate to your environment to try to identify an appropriate storage virtualization solution.

What are your objectives?

Understanding the priorities of your business and how the IT environment can serve to deliver them is critical. If you’ve suffered from application outage or data loss that had a significant effect on your organization, exploring how storage technology can help avoid it happening again is the first place to start. Or if you find the IT budget allocated by the business is simply insufficient to meet the growing data requirements, look for smarter ways of managing your storage resources. Take the time to match these problems to specific capabilities offered by storage virtualization. Ask a number of storage vendors to explain how their solutions can meet your requirements.

Where are you today?

Although the benefits of storage virtualization are enticing and can make a tangible difference to your business, taking stock of where your organization is today in its level of maturity and capability is important. After deciding storage virtualization can help you in your environment, consider your current capacity utilization and how much you can save by reallocating storage resources. Do the vendors you are talking to support the storage components you already have in your environment? How mature are your storage management processes and do they need to be addressed first? Do you have the in-house resources and skills to undertake such an implementation, or will you need to hire contractors or rely on your storage vendor? Do you have the right people to manage the environment after it’s all in place?

Price versus cost

Implementing storage virtualization needs to provide some measurable benefits to the organization; otherwise, let's face it, why would you do it in the first place? Ultimately, stepping into any storage virtualization solution will require an investment of time, people, and money. You will need to go through a discovery phase, assign people to the transition, and possibly buy some equipment and/or software. However, remember that the price of all this does not equal the cost. Implemented correctly, the right storage virtualization system, versus doing nothing, will realize a return from the investment in efficiency savings. It is therefore important to project what these savings will be so that you can evaluate which solution offers the best value. Not all vendors' solutions will be priced the same and some will offer greater capabilities that lead to greater savings, so use your organization's objectives to prioritize the capabilities that are important. In Chapter 3, we offer an approach for doing this.

Matching growth to scalability

Not all organizations grow their storage at the same rate, so you need to evaluate solutions that fit your growth profile. Although all storage vendors like to talk about the massive scalability and performance their solutions provide, you need to cut through the "marchitecture" and only evaluate solutions that best suit your environment. Understand your real capacity utilization, your actual application growth (not storage allocation growth), and shortlist the capabilities you believe will make a difference, perhaps thin provisioning and non-disruptive data migration. Insist on keeping to this scope.

Can you afford an outage?

Every storage solution requires some disruption to your organization while it is implemented. Some solutions require a greater outage to your environment than others, depending on the amount of storage you have and the number and quality of skills you have to carry it out.



As you evaluate solutions, understanding how much disruption implementing the solution will cause is important. You need to feel comfortable that the storage or service vendor can minimize this, not only through the technology itself, but through previous experience.

Chapter 3

The Economics of Storage Virtualization

In This Chapter

- ✓ Measuring the economic benefits of storage virtualization
- ✓ Finding the savings in storage
- ✓ Knowing where to start
- ✓ Using Hitachi's Storage Economics model
- ✓ Looking at storage virtualization in action, a case study

When buying a new car, you don't just consider one aspect, like the upfront cost of purchasing the car itself; you take into consideration multiple factors that determine the overall value for money you are likely to get. You evaluate such things as how much mileage you're likely to get from a tank of gas, what the servicing and insurance costs are, and what safety features are built in to protect you and your precious passengers.

The same is true for storage. If the cost of the capacity were the only factor to consider, then the answer would be simple: Buy the disk from the guy who gave you the cheapest price per terabyte!

However, the impact of continuous growth in storage capacity on the IT budget can be complex, involving the interplay of a diverse set of costs. If all elements in the storage economics equation aren't taken into account, the real cost of a purchasing decision can be decidedly different than intended.

In this chapter, we first look into these elements and just how you can measure and evaluate them. Later in the chapter, we explain Hitachi's unique Storage Economics model, which

provides a methodology for measuring the benefits, and introduce a case study to see how one organization realized significant savings for its business.

How Do You Measure the Economic Benefits?

Identifying cost categories that are important to the organization is the first step in assessing the impact of a proposed storage purchase. We detail these cost categories in the section “Where Do You Start?” later in this chapter. Once identified, costs must be quantified using concise, deterministic methods — such as return on investment (ROI), total cost of ownership (TCO), and return on assets (ROA), which we examine in this section. Other methods of evaluation include time values of money saved, internal rates of return (IRR), and net present value (NPV) of future savings. These metrics allow differences between new and old storage infrastructures to be impartially assessed.

Storage virtualization plays an important part in the storage economics equation as it is an enabling technology to accelerate efficiency and productivity improvements that justify a business case for change.

ROI and TCO are two important mechanisms used to determine the economic value of a proposed purchasing decision. Although the two are frequently treated as interchangeable in vendor marketing literature, ROI and TCO have very different purposes.

ROI is effective when challenging the status quo with a proposition to replace an existing solution. Frequently, IT planners will have an idea that long-term operating expense savings can be realized by making an investment, and the ROI analysis provides the mechanism for fleshing out the financial pros and cons of the decision. The analysis illustrates how much is to be invested, how quickly the investment is to be recouped, and what net savings are to be expected. Figure 3-1 shows a ROI analysis over three years.

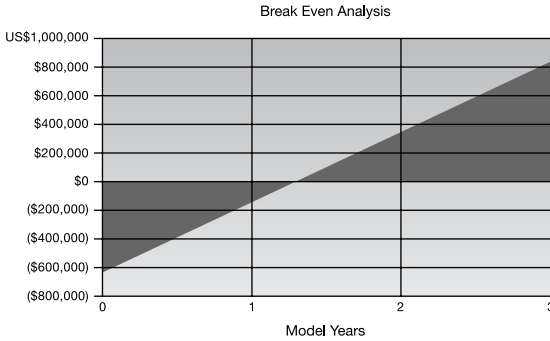


Figure 3-1: This return on investment (ROI) over three years shows a break-even point a little into the second year.

TCO analysis is used when the financial benefits of two or more proposed solutions, or a solution against the status quo, must be assessed. The analysis places the total lifetime operating and purchasing costs of the assets side by side for comparison (see Figure 3-2). IT purchasing best practices, documented by independent analysts, frequently cite the importance of including a TCO analysis in any competitive bid situation. The cost categories chosen for comparison are those deemed most significant, and these frequently vary from one organization to the next.

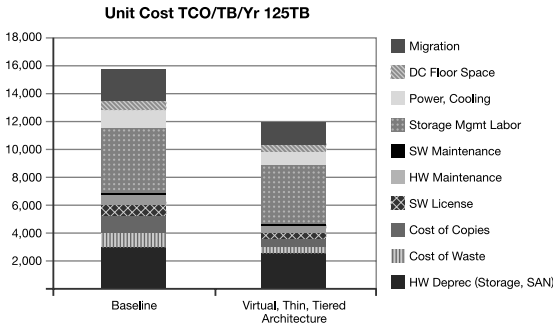


Figure 3-2: Total cost of ownership (TCO), comparing the status quo with an alternative.

However, ROA is a broader and more accurate measure of a company's ability to make profit from its invested capital. ROA is the ratio of net income to total assets. Many CxOs often regard IT storage as having a relatively poor ROA due to the historically low levels of capacity utilization caused by over-provisioning, data duplication, RAID (redundant array of independent disks) overhead, and stranded capacity. With the advent of storage virtualization, and the ability to repurpose and reclaim legacy assets, thereby increasing their value, organizations can help realize a faster return to the business, accelerate the savings, and reduce the TCO.

So Where Are the Savings in Storage?

For several years, Hitachi Data Systems has conducted exercises for hundreds of organizations around the world, assessing the economic impact of storage on their business. This work has led to an astounding data point. On average, for every 12 terabytes (TB) of usable disk capacity within the storage infrastructure, there is potential for \$1-million savings in net operating expense.



This single data point, which is often sufficient to justify further investigation and economic analysis, is reached through savings across the business environment. Here are the average savings in each category:

- ✓ Reduced waste, 25 percent
- ✓ Reduced outages, 20 percent
- ✓ Reduction in labor and management, 15 percent
- ✓ Reduced maintenance fees, 15 percent
- ✓ Environmental savings, 10 percent
- ✓ Miscellaneous operating efficiency improvements, 5 percent
- ✓ Other savings, 10 percent

Where Do You Start?

The first step in assessing the impact of a storage system on the IT budget in your business is to determine what expense categories are real and defensible to the organization. Operating expenses (OPEX) fall into two primary categories: hard costs, the tangible and visible expenses that find their way into budgets and financial reports; and soft costs, those less tangible but equally important measures, such as performance and availability. Each organization will have its own idea of which costs are hard and which are soft.



Primary areas of investigation should include:

- ✓ **Data protection.** Includes the cost of the organization's backup and business continuity infrastructure, which may include tape, replication, mirror sites, and all related planning and operational activity.
- ✓ **Maintenance.** The cost of hardware and software maintenance is linked to the capacity deployed, whether in use or not. It may increase as hardware warranties and bundled service deals expire.
- ✓ **Staff costs.** The labor cost of storage administration and provisioning can be significant, particularly in complex heterogeneous environments that may have multiple management interfaces and distributed architectures.
- ✓ **Environment costs.** Power, cooling, and floorspace costs are increasingly important as companies strive to meet efficiency targets that are often driven as much by a corporate and social responsibility agenda as the need to save money. A significant number of companies are reaching or have reached absolute limits on their data center growth and power consumption.
- ✓ **Outage costs.** These costs will vary widely depending on the organization type, industry sector, and application concerned. A realistic estimate of outage costs is a key element in planning the company's business continuity arrangements.
- ✓ **Cost of growth.** This will depend on the storage tiering structure, since each tier is typically subject to different growth rates, capital expenditure (CAPEX), and hardware price erosion. For example, lower tiers tend to grow more quickly than higher tiers and the price of low-cost SATA (serial advanced technology attachment) hard drives

tends to fall faster than more reliable fibre channel (FC). Redistributing data between tiers is also a cost.

- ✔ **Information governance.** An increasingly important element in the total cost of storage, with obligations for companies to hold data for longer periods of time with a demonstrable search and recovery capability.
- ✔ **Migration.** Bringing new capacity online, taking old arrays out of service, or moving applications between platforms can involve migration projects that may be time consuming and expensive, particularly in non-virtualized environments.
- ✔ **Performance.** Extra costs are incurred by the need to deploy storage systems of sufficient performance and scalability to meet the requirements of the service level agreements (SLAs) with the IT users, and thus defer the need for further investment as user capacity needs and productivity grow. This may also mean that I/O performance upgrades are added at the same time as storage capacity upgrades.

Hitachi's Storage Economics Model

Hitachi Data Systems, through hundreds of TCO and ROI analysis studies, has evolved an approach to storage economics and developed a set of tools and methodologies that can be applied to a broad range of customer types across different industry sectors.

The basis of this model compares the existing infrastructure (business as usual) with a new solution that includes modern storage innovations such as multiple storage tiers, storage virtualization, and thin provisioning technologies (explained further in Chapter 5).

This helps storage managers and CxOs gain a deeper understanding of the true operational expenses in their storage infrastructure and establish a roadmap for future cost reductions. This thorough analysis combines a set of technical reports with structured customer interviews that provide an assessment of the storage infrastructure's operations and costs in light of the individual business objectives. The methodology follows three key stages:

- 1. Establish a TCO baseline.** What does it cost the organization to own and operate 1TB of storage per year? The answer is derived from an understanding of the unique set of costs that apply to that storage customer, identified in up to 33 different cost areas ranging from power and cooling to labor charges and outage costs.
- 2. Map the costs to a set of strategic and tactical cost-saving initiatives.** What cost-saving initiatives are needed? These may be strategic or tactical initiatives, such as consolidation, tiered storage, virtualization, thin provisioning, chargeback, or implementation of management tools, and will vary depending on the customer's environment. The user maps the costs identified in the first step to the cost-saving initiatives, so that the multiple relationships between costs and investment can be established. This process also helps to identify the areas of technology investment that are the highest priorities.
- 3. Compile a cost reduction roadmap.** What will be the impact on the TCO baseline over the short and medium terms? With a rollout plan for each of the cost-saving initiatives, a roadmap can be built that shows the impact on each of the key cost areas year by year.

By following this three-stage approach, a tactical plan for addressing the key cost areas and driving long-term TCO reduction can be scheduled and executed.



Four principles of storage economics

Hitachi has developed an online tutorial outlining four principles to reduce the total cost of ownership (TCO) with storage virtualization.

Go to www.hds.com/economize.

Case study: Indosat's storage solution

Indonesian Satellite Corporation Tbk (Indosat), one of Indonesia's largest telecommunications providers, was formed with the merger of three separate entities. The merger transformed the company into a cellular-focused, full-network service provider servicing over 12 million customers, and presented a major challenge.

Working with three different sets of storage systems threatened to take a toll on Indosat's IT infrastructure operations and maintenance. It was clear that direct attached storage for each server was no longer an option as it would not meet Indosat's requirements for scalability and flexibility.

In addition to this challenge, a rapidly growing customer base was straining the capacity of the existing systems. A disaster recovery facility was also needed to protect critical application data in the event of systems failures. Consolidation of the IT infrastructures, as well as the applications, became essential.

Indosat identified Hitachi Data Systems as the solutions provider of choice to quickly take control of the situation. By drawing on a combination of skills in consolidating storage resources and leveraging

the unique virtualization capabilities of the Universal Storage Platform, Hitachi Data Systems designed a solution that integrated the three storage infrastructures seamlessly into one environment while leveraging existing investments. Indosat's applications were classified into relevant tiers that matched the value of the data to the cost of the storage, ensuring the organization was confident they had a sustainable way of optimizing storage resources that raised efficiency and accountability.

Virtualization paved the way for consolidation and allowed Indosat to make the best use of their existing storage, effectively protecting the bottom line. Virtualization also opened the door to a host of other benefits, such as reduced management of provisioning, ease of monitoring to maintain service levels, and flexibility to conduct data migration and replication without increased risk.

The results speak volumes of the project success: 10 to 20 percent savings in human resources and 40 percent savings in hardware investments. Needless to say, there has been no looking back for Indosat.

Chapter 4

Making Storage Virtualization Real

.....

In This Chapter

- ✓ Understanding key concepts
 - ✓ Planning to virtualize your storage through objectives and analysis
 - ✓ Implementing your virtualized storage solutions
 - ✓ Migrating your data
-

Storage virtualization can help address many shortcomings of traditional storage infrastructures, enhance the storage investment, and respond to real business problems. Walk into almost any data center with multiple storage platforms from different vendors and you're likely to find low storage utilization, multiple management tools with different interfaces, and disparate storage management practices for each vendor's storage. Storage virtualization solutions help you to simplify the storage infrastructure and dramatically improve storage utilization by managing all storage assets as a single pool via a single management interface.

Data migrations are part of any ongoing IT operation and are especially intensive during technology refresh cycles. Migration is highly prone to human error and typically needs to be done in very small windows of opportunity. A virtualized storage environment enables an intelligent and dynamic tiered storage solution that allows seamless data migration across storage tiers without impacting users or applications. Intelligent tiered storage solutions help you to simplify infrastructure, ensure quality of service, reduce risk, and align the right storage tier to the right application, thereby reducing capital (CAPEX) and operating (OPEX) expenditure.

In this chapter, we deal with the actual implementation of storage virtualization, but first we outline some of the key concepts used in storage virtualization.

Key Concepts in Storage Virtualization

Any technical field develops its own terminology based on the type of concepts involved, and storage virtualization is no different.



Here are some of the key ideas to get your head around:

- ✓ **Address space remapping:** Virtualization of storage helps achieve location independence by abstracting the physical location of the data. The virtualization system presents to the user a logical space for data storage and handles the process of *mapping* it (making it visible) to the physical location. The actual form of the mapping will depend on the chosen implementation.

In a block-based storage environment, a single block of information is addressed using a logical unit number (LUN); an offset within that LUN is known as a logical block address (LBA). The address space mapping is between a *logical disk*, usually referred to as a virtual disk, and a *logical unit*, presented by one or more storage controllers. The LUN itself may also be a product of virtualization in a different layer.

- ✓ **Meta-data:** The virtualization software or device is responsible for maintaining a consistent view of all the mapping information for the virtualized storage. This information is usually called *meta-data* and is stored as a mapping table. The address space may be limited by the capacity needed to maintain the mapping table. This capacity is directly influenced by the *granularity* (the level of detail) of the mapping information.
- ✓ **I/O redirection:** The virtualization software or device uses the meta-data to redirect I/O (input/output) requests. The software or device receives an incoming I/O request containing information about the location of the data in terms of the logical disk (virtual disk), and translates this into a new I/O request to the physical disk location.

- ✓ **Pooling:** The physical storage resources are aggregated into storage pools, from which the logical storage is created. More storage systems, which may be heterogeneous in nature, can be added as and when needed, and the virtual storage space will scale up by the same amount. This process is fully transparent to the applications using the storage infrastructure.
- ✓ **Disk management:** The software or device providing storage virtualization becomes a common disk manager in the virtualized environment. Logical disks are created by the virtualization software or device and are mapped to the required host or server, thus providing a common place or way for managing all volumes in the environment.

Planning for the Virtualized Storage Environment

First and foremost, it is important to understand that storage virtualization is not something you buy off the shelf at your friendly IT supermarket, plug in, and off you go. Storage virtualization is a process and, just like any good vacation, the fun is in the journey as much as the destination. By starting small, in one part of your environment, you'll get some benefits. As you extend storage virtualization to other applications, steadily the benefits compound, and then you can introduce more advanced features like migration and thin provisioning to really accelerate the savings.

After deciding storage virtualization can make a difference in your environment, planning for its introduction is critical. Much like when you take a driving vacation, you check that you have enough oil in the car, check the tire pressure, use maps to know where you want to go, devise a route, and decide where you'll stop for gas.

Every organization is unique, which means that its IT requirements and priorities are different. Establishing your objectives and priorities up front is important, because the success of storage virtualization is largely driven by an organization's understanding of its needs. You need to establish how critical an application is to the business, as well as its availability and performance. Once the objectives and priorities are understood, the storage architecture, along with storage management standards, policies, and operational procedures,

can be devised to meet them. Especially important here is to identify how storage virtualization can enable you to fulfill those objectives and priorities.

If the primary objective of the IT organization is to improve the storage utilization and save on the CAPEX and OPEX by deferring the storage purchases and decreasing the administration burden, you need to understand the real scenario in the infrastructure. What is the current utilization level and how much improvement does it need? Where are the storage devices and how are they connected to the servers, applications, and so on? A good number of SRM (storage resource management) tools work in either a homogeneous or multi-vendor environment and report on this. Element managers from respective vendors also give this information for reference.

If the objective is to reduce the risk and downtime, and improve the service levels, then you need to do overall performance analysis of the infrastructure and set the baseline for monitoring the difference. A good number of tools are available to do this as well. Unplanned downtime or unexpected outages of the storage can impact business very badly, so understanding this well is critical.

Once the objectives and priorities are set, you're ready to choose the right method of virtualization (Chapter 2 discusses different types of storage virtualization to choose from). At this time, a good strategy is to create a "storage services catalog" and classify the data storage as required by applications for the "right" type of storage based on the performance and capacity needs. We discuss this in more detail in the section "Migrating Data to a Virtualized Storage Environment" later in this chapter.

Implementing Storage Virtualization

After you've identified the need for consolidation of your storage environment and planned your storage virtualization, the next step is to actually make the required logical and physical changes to the storage environments.



Consolidating and virtualizing storage devices may create performance issues for the critical applications. Following best practices to prevent any bottlenecks that may arise and using the right tools to monitor the environment is important.

Partitioning, the subdivision of components such as disks, cache memory, and ports to configure storage arrays, can assist in ensuring that these critical resources aren't overloaded by one application, which may adversely affect another application.

Partitioning of storage hardware

Virtual partitioning enables logical partitioning of the physical ports, cache memory, and disks on the virtualized storage platform. You can then improve security by “restricting access” so that the administrator of one partition won't be able to access or interfere with priorities or data in another partition. Virtual partitioning also allows you to optimize application performance within a partition. By dedicating a fixed amount of cache memory to a partition, you can guarantee that the application is getting the required resources and won't be affected by any other application, host, or user.

Partitioning also enables you to take a departmentalized view of storage (that is, a view of the common storage pool that is virtualized). This ability offers the benefits of centralized management and control, alongside decentralized departmental control and accountability.

Internal and external storage configurations

For controller-based storage virtualization platforms, zones need to be created in the storage area network (SAN) that include the front-end storage ports of the virtualized storage platform and the ports of the external storage platforms being virtualized.



Tasks involved in configuring virtualized storage environment are:

- ✓ **Configuration of external ports.** Virtualized storage platforms can support several types of ports. A *port* is a connection point where the storage array is connected either to another storage device or to the SAN switch. Target ports (or *front-end ports*) are used by hosts or SAN switches to connect to the virtualized storage platform device. External ports (or *back-end ports*) are used to connect to external storage.

- ✔ **Discovery of externally attached storage.** Using the defined back-end ports on the virtualized storage platform, you need to recognize any externally attached storage and its LUNs.
- ✔ **Discovery of LUNs configured within externally attached storage.** Once storage is discovered, LUNs can be mapped and managed within the virtualized storage platform's software.
- ✔ **Creation of external groups of storage (eGroups) and external LUNs (eLUNs).** This creates LUNs within the virtualized storage platform, utilizing the defined external storage.
- ✔ **Mapping of eLUNs to hosts.** The new eLUNs can be allocated to host systems within the environment, just as for the internal LUNs.

Figure 4-1 illustrates how you can configure a virtualized system using a virtualized storage platform, with two external storage arrays.

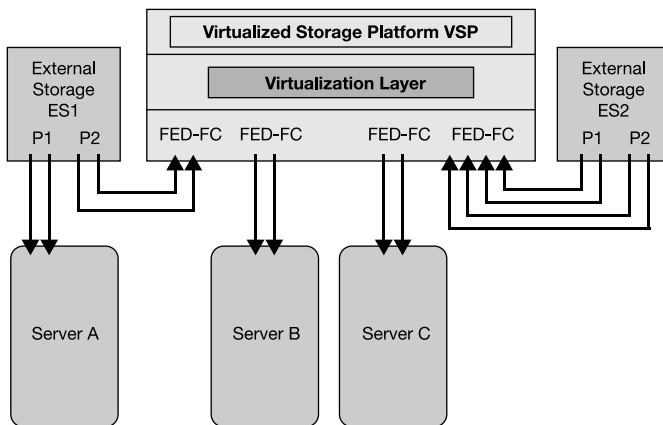


Figure 4-1: This example of a virtualized storage system uses two external storage arrays.



The sample system shown in Figure 4-1 features the following characteristics:

- ✓ Two storage systems (external), ES1 and ES2, are virtualized behind the virtualized storage platform (VSP).
- ✓ Three servers, Server A, Server B, and Server C, access storage from these storage systems.
- ✓ Storage LUNs from ES1 are migrated onto the VSP's internal disks (which are more likely to be high-performance disks).
- ✓ Server B and Server C access storage via the VSP, either internally (for high-performance storage needs) or through storage from ES1 or ES2, which is virtualized.
- ✓ Server A continues to use those LUNs from ES1 storage that are not virtualized behind the VSP.
- ✓ New LUNs will be created within the VSP from its local disks if Server B and Server C need additional high-performance storage to support new production applications. Alternatively, the storage virtualized from ES1 and ES2 can be provisioned to Server B and Server C.

Virtual partitioning can create separate areas and isolate the effects of the various workloads from one another. For example, you can define individual storage partitions to isolate management of one server's allocated storage from another server's allocated storage. You can also create cache partitions to assign specific cache amounts to the VSP's internal storage, ES1, and/or ES2.

Migrating Data to a Virtualized Storage Environment

After the existing storage is virtualized behind the controller, you can plan for migration of data from older storage to new fast-performing storage devices.



Some applications may need very high-performance, highly redundant storage disks with faster rotation time, whereas others can use storage from lower performance disks. An application's needs may change over time; therefore the ability to migrate from one type of disk to another on demand is extremely valuable. Ideally you should perform these migrations without stopping the applications.

Creating tiers of storage is also a good practice, selecting the right type of disks or LUNs to form a tier for the purpose of migrating or provisioning. However, locating disks with the right characteristics from the hundreds and thousands of disks in the storage environment can be a big challenge.

Classifying the data types

Classification of data is the assignment of different categories of data to different types of storage media in order to reduce total storage cost. Categories may be based on levels of protection needed, performance requirements, frequency of use, and other considerations. Since assigning data to particular media may be an ongoing and complex activity, some vendors provide software for automatically managing the migration process based on a predefined policy.



Valuable or frequently accessed data, for example, is kept on high-performance fibre channel (FC) disks, whereas less valuable data is moved to less expensive near-line storage, such as a serial attached SCSI (SAS) or serial advanced technology attachment (SATA) disk, and infrequently accessed data can be migrated to a high-volume SATA disk or tape. Ideally, utilizing tiered storage can save money, while easing the access demands on any single storage tier.

Creating storage tiers

Tiered storage is a strategy deliberately undertaken by an IT organization to distribute different classes of data across differently valued storage elements. These tiered systems are distinguished by differences in cost, performance, reliability, and feature sets.



Tiered storage may also include different classes of disks within the same storage system, such as 15k RPM (rotations per minute) versus 10k RPM disks, or FC and SATA disks (on mid-range systems). A migration of a LUN copies the contents of the source LUN to a target LUN. At the point of complete synchronization between the two LUNs, the LUN identifiers are swapped. Any host accessing the LUN will not see a change since the LUN identifier did not change.



Although storing everything the same way may appear easier — on the same class of storage system with the same protection levels — doing so creates unnecessary costs. The task of classifying data is somewhat daunting and requires time that most companies don't have. A tiered storage approach helps to move data based on volume groups and their corresponding application requirements, which relieves companies from the arduous task of micro-managing the classification of all data.

Uninterrupted data migration

Data migrations are extremely common in storage operations. Server and storage equipment eventually needs replacement (typically called *technology refresh*) due to maintenance contracts expiring, lease renewals, or no longer being supported by vendors. Equipment sometimes needs to be relocated to another data center or consolidated with other equipment. Other reasons to perform data migrations are to move infrequently accessed data to lower cost storage systems or promote data to higher performance storage, as business dynamics dictate.

Traditional methods of data migration are inherently disruptive because the process, involving multiple people in the IT organization (application, system, network, and storage administrators), can take a great deal of time and resources, and poses risks if something goes wrong. This is why many organizations are demanding less intensive ways of managing data migrations, especially among different heterogeneous storage tiers.

By implementing storage virtualization, heterogeneous migration of data between one vendor's storage to another's is possible. Defining groups of volumes according to how they relate to the application and applying policies by way of highly customizable storage tier definitions enables the movement of data without interrupting running applications.

Figure 4-2 illustrates the data migrated online from an active volume from the external storage, ES1, to the tier-one high-performance internal storage on the virtualized storage platform. The example shows a man watching a streaming video played on virtualized storage attached to the server. With dynamic tiered storage, migration can be executed without any disruption to the video playing, since it is completely transparent to the server.

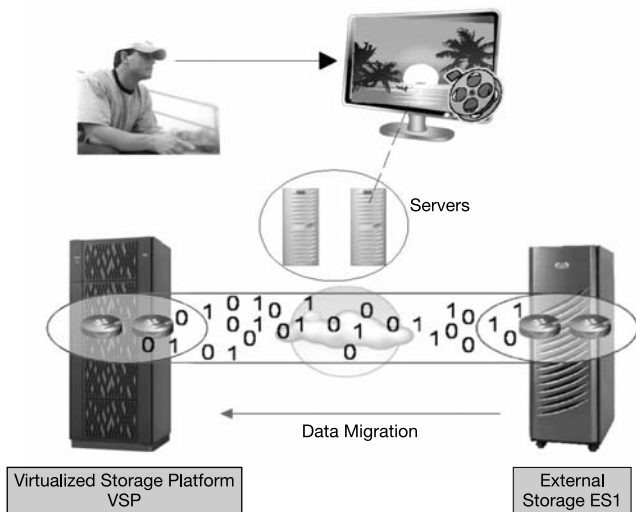


Figure 4-2: In this example, the streaming video is uninterrupted during data migration.

Chapter 5

Getting the Best from Storage Virtualization

.....

In This Chapter

- ✓ Looking at internal storage virtualization through thin provisioning
 - ✓ Understanding management challenges and solutions
-

Rapid storage growth, with its associated requirements for provisioning and management of storage assets, has become a major challenge for information technology. Traditional static provisioning often results in over-provisioning with significant amounts of unused capacity, which negatively impacts storage total cost of ownership (TCO). In this chapter, we look at how internal storage virtualization can improve both application performance and the utilization of storage capacity. We also delve into the management challenges faced by organizations when they implement virtualized storage systems, as well as the tools and functions required to meet them.

Optimizing Performance and Improving Storage Utilization

Internal storage virtualization comes in different forms and capabilities within a storage system. For example, the ability to abstract all of the disks in large logical pools is a core storage virtualization technology that provides significant value. Logical storage pools make it easier to manage and provision your disk capacity. Additionally, they can provide a huge performance boost by striping data across dozens, even hundreds, of disk drives, making all of the disk actuators work in concert to perform read/write operations. Data *striping* is a process that takes multiple physical disks, breaks up the files into even pieces, and spreads them across multiple storage devices.

These materials are the copyright of Wiley Publishing, Inc. and any dissemination, distribution, or unauthorized use is strictly prohibited.

Another valuable internal storage virtualization technology is *thin provisioning*, also known as *dynamic provisioning*. This technology allows for greater capacity utilization, brings more applications online without having to buy capacity in advance, and simplifies the storage provisioning management process. In this section, we look at various aspects of thin provisioning.

Understanding thin, or dynamic, provisioning

Thin, or dynamic, provisioning can be used to present a virtual pool of shared capacity that is larger than the actual amount of physical storage available. It enables system administrators to deliver capacity on demand from a common pool of storage. What differentiates it in the virtualized storage environment is the ability to dynamically provision from both internal and external storage, thus allowing advanced virtualization capabilities on older or lower tier arrays.



The difference between traditional provisioning and dynamic, or thin, provisioning is shown in Figure 5-1. In the case of a server that requires 400 gigabytes (GB) today but is projected to grow to 2 terabytes (TB), traditional methods would require the full 2TB to be allocated today to avoid lengthy downtime for volume migration or expansion as the data set grows (as shown on the left). Here the free capacity is not available to re-assign to any other servers since it is already allocated. Dynamic provisioning allows a 2TB volume to be presented to the server with only 400GB of physical capacity actually used (as shown on the right). In this case, 1.6TB of free capacity is actually available to other servers since it is not allocated. If and when the amount of capacity approaches 400GB and exceeds defined thresholds, additional physical capacity can be added online with no impact to users or applications.

Creating a dynamic provisioning pool

Dynamic provisioning is implemented using a concept called a *dynamic provisioning pool (DP pool)*. A large number of drives can be part of a DP pool, and the storage controller automatically stripes data across all available disk drives in the pool. Implementing a pool over a large number of disk drives is also known as *wide striping*.

Figure 5-1 (on the right-hand side) demonstrates dynamic or thin provisioning storage pools and thin provisioned volumes. The host servers see the virtual capacity, represented by the dotted line disks, whereas the actual capacity is in the DP pool, represented by the stripes just above the LUN (logical unit number) parity group.

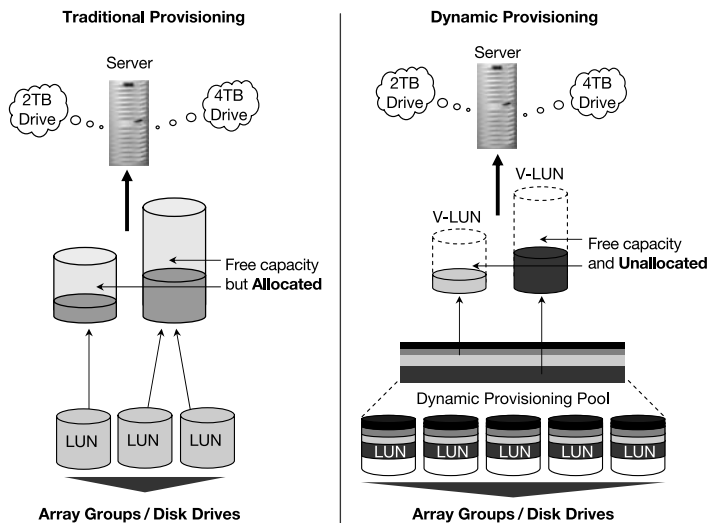


Figure 5-1: Comparing traditional provisioning with dynamic provisioning.

Benefits of dynamic provisioning

Dynamic provisioning introduces many benefits, such as:

- ✓ Closely coupling storage purchases to the application's consumption.
- ✓ Repurposing unused storage that is already allocated, using the storage only when the host (as per the application's demand) writes the data on it.
- ✓ Simplifying storage management tasks by managing capacity from a central pool instead of server by server.
- ✓ Reducing capital costs, as over-provisioning is reduced or eliminated.
- ✓ Improving performance, as data sets are striped over a large number of drives.

When combined with storage virtualization, the benefits of dynamic provisioning can be extended to the external virtualized storage to save space and to improve performance through wide striping.

Space saving

Local and remote volume copy services that are “thin provisioning-aware” consume less capacity. With traditional provisioning, if you have a 2TB volume but only 100GB is actual data, all of the 2TB will still be copied. In some cases, users will create two or more local copies of their primary volumes. With traditional provisioning this would consume lots of capacity by copying empty blocks; thin or dynamic provisioning eliminates this.

One technique in dynamic provisioning allows the physical capacity to be reclaimed if not used by an application, even if the capacity is allocated to the application. Using dynamic tiered storage data migration technology, data on a physical “thick” volume can be migrated to a logical “thin” volume, and unused allocated space can be reclaimed. This is called *zero page reclaim*.

Thin provisioning can enable more applications and servers per storage system, providing greater levels of consolidation. Since capacity doesn’t have to be dedicated and fixed on a per-volume basis with thin provisioning, customers can create more volumes. As a result, more applications and servers can be attached to a single storage system.

Due to thin provisioning’s capacity efficiency, it requires fewer disks, which results in lower power and cooling costs. Often this is essential — for data centers that are running out of power, for example.

Wide striping for performance improvements

Mechanically spinning hard drives are the slowest component in the storage performance chain. The implementation of additional drives behind each application volume enables more work to be done in parallel. This is particularly true for massively

interactive applications, including online transaction processing and email or messaging systems.

This phenomenon is well known by database and storage administrators, who have been using host-based volume management software for years to make multiple fibre channel (FC) LUNs appear to an application as a single volume. Thin- or dynamic-provisioned volumes defined over a wide-striped pool of drives can be used to achieve the same performance benefits.



A major benefit of large logical pools of capacity is the ability to stripe data across a large number of disk drives. This improves performance significantly, since you can have dozens or hundreds or even thousands of disk drives conducting I/O (input/output) operations simultaneously.

Additionally, wide striping significantly reduces performance management tasks. A great deal of time is often needed to analyze, tune, and retune performance with traditional provisioning methods. Wide striping minimizes and potentially eliminates the need to tune and analyze, since it is perpetually optimized for performance.

Management Challenges in the Virtualized Storage Environment

Many environments face ever-growing complexity in the underlying infrastructure needed to support their core applications. At the same time, IT organizations are being asked to do more with lower budgets and the same or decreasing levels of staff. Therefore, investing in a simplified storage management infrastructure that supports the goal of doing more with less is more critical than ever.

Examples of the increased complexity found in today's enterprise data center include:

- ✓ Support for multiple host operating systems
- ✓ Data storage for multiple terabytes (TB) or petabytes (PB) of application and file data
- ✓ Multiple storage systems in the environment, often from different vendors

- ✔ Storage, often existing as isolated islands, or “silos,” each with separate management tools
- ✔ Application downtime that is unacceptable for routine tasks (protection, adding capacity, management) on key production systems
- ✔ Growing concern regarding data center “hot spots” and constraints on physical space as capacity increases



TIP

Storage consolidation with storage area networks (SANs) has increased utilization but introduced the likelihood of performance problems. Applications are more exposed to performance issues than any kind of component or data center disaster, so understanding their workloads is very important. Storage infrastructure is more complex than ever (SAN, cache, ports, virtualization, thin provisioning). *Baselining* (analyzing current performance to provide a baseline for comparison with other systems) and monitoring can be difficult without the right tools, but it is critical for implementing storage virtualization successfully.

Ideally, there should be a minimum set of tools that can manage and monitor the virtualized storage environment. However, in reality, IT administrators need to use multiple sets of tools to do point-in-time management for its various components.



REMEMBER

Key points to consider in selecting the right management software or tools are that each tool should:

- ✔ Provide end-to-end visibility of the complete storage infrastructure so the physical connectivity between applications, server, SAN, the virtualization storage platform, and external storage is understood.
- ✔ Collect detailed storage configuration attributes, such as model, type, status, ports, mappings, RAID (redundant array of independent disks) groups, and LUNs.
- ✔ Have the ability to search and filter storage characteristics (such as size, protection, and performance) to create and manage storage tier definitions and migration groups.

- ✔ Monitor end-to-end performance of not only application, server, SAN, and storage, but also virtualization attributes such as external storage ports, logical partitions, and cache.
- ✔ Provide detailed monitoring and reporting of thin provisioning pools (virtual and physical utilization), as well as the ability to issue alerts when resource thresholds are met.
- ✔ Monitor, prioritize, and even cancel migration and replication tasks to ensure complete control.

Dealing with the challenges

Storage virtualization decreases complexity, allowing one administrator to manage three to ten times more storage.



With Hitachi Data Systems' Universal Storage Platform series, storage virtualization can simplify the management of otherwise complex storage infrastructures in many ways:

- ✔ **Connectivity and interoperability.** Since the Universal Storage Platform mimics a server host connection to external storage, interconnecting heterogeneous storage assets does not require special certification effort or changes to microcode. Furthermore, since it uses a standard connectivity protocol (SCSI, or small computer system interface), it supports more storage systems on the market.
- ✔ **Unified management.** Using a common management tool set, the Hitachi Storage Command Suite, multiple software products can be consolidated. In addition, the burden of management is significantly reduced by using an interface for provisioning, management, reporting, and protection.
- ✔ **Business agility:** Utilizing the power of an intelligently switched architecture, increased levels of scalability and performance are available to applications. The platform also enables non-disruptive data migration to transition and transform the storage infrastructure to meet changing business requirements.

Managing performance and capacity

Earlier in this chapter, we discuss the new challenges IT organizations potentially face when introducing new technologies such as virtualization and IT consolidation, as well as the management tools and functions needed to deal with those challenges. Performance and capacity management of the virtualized environment is also essential.



With the ability to partition storage resources into multiple virtual machines, each with its own dedicated external capacity, cache, and ports, you can minimize the chances of a rogue application consuming all the resources and impacting other applications.

Furthermore, since each virtual storage machine looks like a separate, isolated storage device, IT teams can associate specific business applications with selected departments and business units, then charge back accordingly. These capabilities allow organizations to create a utility model to deliver metered storage resources to select business entities. Subsequent management of virtual storage machines can be addressed centrally or by local administrators.

Measuring key metrics when applications and infrastructure are running well is an important task for maintaining service level agreements (SLAs). These baselines serve as a reference point to continually assess the infrastructure against. Any resources (CPU, bandwidth, storage) that have the potential to impact application performance, and therefore the service level, can be easily identified. Once recognized, thresholds can be set and alerts defined to ensure that the IT organization can act upon them before a problem arises.

Chapter 6

Ten Best Practices for Deploying Storage Virtualization



In This Chapter

- ✓ Putting in place a process for establishing physical configurations
- ✓ Choosing the right way to implement data migrations and tiered storage
- ✓ Getting the most from thin or dynamic provisioning
- ✓ Discovering scenarios in which storage virtualization can improve IT efficiencies



In this book, we give you a pretty good primer on storage virtualization — the concepts, benefits, and different types. We examine the value of storage economics to validate the savings that storage virtualization can bring. We also look at how to plan for different techniques and methods used in storage virtualization, such as data migration, tiered storage, and thin, or dynamic, provisioning.

In this chapter, we discuss a design methodology and best practices for developing and deploying the storage virtualization technologies.

#1: Choose the Right Storage for the Host's Mapping

In order to implement storage virtualization for physical configurations, as a first step, you need to gather information

on all host systems and disk configurations. Then allocate the disks to the hosts from the storage subsystems by discovering them from the host bus adapter (HBA). To ensure that the target logical unit numbers (LUNs) are available on the storage area network (SAN), the best practice is to use an HBA tool (such as HiSat from Hitachi Data Systems) to locate them. This way you know that the mapping is done properly. In certain cases, this process may require a reboot.



Having multiple paths for storage resources is always recommended for load balancing and *failover* (automatically switching to a different redundant system in the event of failure) or redundancy. Increasing the number of paths from the virtualized storage platform to the external storage results in a significant increase in the speed of data transfer. Updating I/O (input/output) activity will have a major impact.

#2: Assign External Storage Configurations

After selecting the right storage for the host's mapping (refer to the previous section), the next step for physical configurations is to assign external storage groups (eGroups). You also need to assign external LUN (eLUN) addresses to the LUNs that have been discovered by the virtualization layer on the virtualized storage platform for the external storage systems.



Disabling cache (cache mode) while assigning eLUNs for data movement purposes is very important. Having cache enabled forces the use of the virtualized storage platform cache for updates. If, at the time of de-virtualization, all segments from cache were not de-staged, data corruption occurs — which can result in serious customer data integrity issues. When creating host storage mapping, be sure to specify the right host modes for the server that will be connected (such as Windows or UNIX).

#3: Format the Virtual Disk

The next step in implementing storage virtualization for physical configurations is to use the virtualized LUNs instead of the native LUNs. To do this, you need to delete the partitions from the native LUN and re-create the partitions using the operating

system command (such as the Solaris format). When virtualizing an existing LUN, a good practice is to run the format command in order to display the partition table of the native volume. Doing this will allow you to re-create the partitions on the new virtual LUN using the right geometry.

After virtualizing the storage, the host may see two disks for the same physical device, one for the virtualized device (target) and the other for the original device (source). Understanding the source device in the controller–target–device (such as c1–t1–d1) combination and correlating it with the target device (say, c2–t2–d2) is critical after the virtualization and volume migration. This will help keep track of the actual storage model and device, and keep the integrity of the data that resides on the disk.



Do not use the Auto Configure option during the formatting as this can modify your disk label.

#4: Choose the Right Disk Attributes for Data Migration

After the physical connections and host storage mappings are defined (refer to the previous sections), you're ready for data migration. This step is about migration of data from a volume in one storage subsystem to another. For the migration to take place, a target volume with the right compatibility must exist. Key criteria are size and emulation mode.



When performing the migration of volumes using a virtualization storage platform that has other workloads running, the best practice is to use a dedicated virtual storage partition. This way the migration will have minor impact on the other workloads.

#5: Migrate the Data

Use the following guidelines to achieve an effective migration process:



- ✓ Do not perform migrations when the virtual storage platform's cache write pending rate is higher than 60 percent.

- ✔ Make sure that the updating I/O rate is below 50 IOPS for the source volume.
- ✔ Do not perform multiple migrations targeting the same RAID (redundant array of independent disks) or array group as this could adversely affect performance.
- ✔ Monitor the RAID group busy rates while selecting the target LUNs for migrations.
- ✔ Do not perform migrations during the following operations:
 - Backups or restores
 - Application upgrades
 - Month-end processing
 - Quarter-end processing
 - Year-end processing
 - Payroll cycle

#6: *Classify the Data*

The design of a virtualized tiered storage system should start from the application's perspective. The business needs and applications drive the storage requirements, which in turn guide tier definition. Most applications can benefit from a mix of storage service levels, using high-performance and less expensive storage depending on the usage and business needs.



Recommended practice is to develop a storage services catalog of tiers with predefined characteristics and then allocate storage to applications, as needed. Figure 6-1 shows a multi-tier model outlining main classes of data tiers — production data, protection data, and archival data. Production data supports your primary applications, whereas protection data stores copies of data for recovery. Archival data is for stale but still important data (possibly on lower cost storage) for compliance reasons. Your individual requirements may call for more or less.



Most customer environments contain a mix of workload and applications within a single system, so when you design the storage system make sure that I/O load or activities (such as formatting, shredding, replicating, or parity rebuilding) on a tier do not disrupt or degrade other tiers. In particular, don't allow a lower tier load to affect higher tier performance.

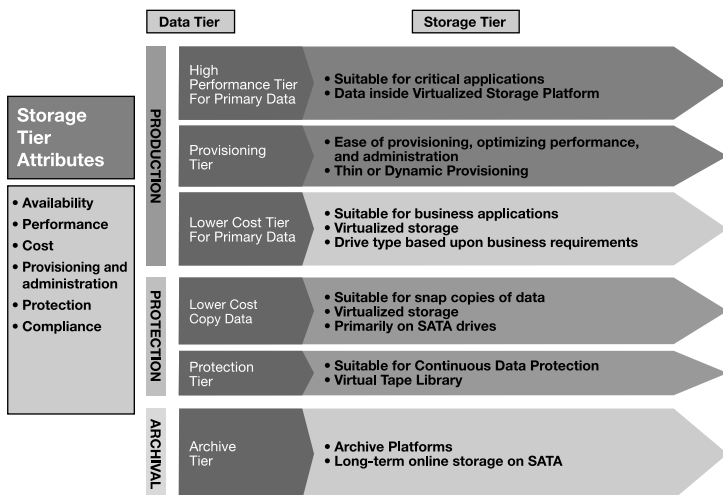


Figure 6-1: Types and attributes of data and storage tiers.

#7: Understand the Workload or I/O Profile

In looking at performance requirements for designing or selecting a storage tier, understanding the application's use of that tier's storage is also important. The mix of read/write characteristics is a major determining factor in the storage performance it will see.



Ask yourself:

- ✓ How much is a sequential read?
- ✓ How much is a random read?
- ✓ How much is a sequential write?
- ✓ How much is a random write?



Consider an I/O's data path inside the virtual storage platform when answering the preceding questions:

- ✓ Random read hits are serviced at fastest speed by cache and do not reach the back end.

- ✔ Random read misses go through cache unaltered and go straight to the appropriate back-end disk drive. This is the only type of I/O operation where the host always sees the performance of the back-end disk drive.
- ✔ The host sees random writes completed at the fastest speed; it only sees delay if too many pending writes build up. Each host random write is transformed when going through cache into a multiple I/O pattern that depends on the RAID type.
- ✔ Host sequential I/O reads and writes are done at fastest speed. The cache acts like a holding tank and the storage system back end puts (or removes) “back-end buckets” of data into (or out of) the tank to keep the tank at an appropriate level.

High-performance drive type and RAID configurations will have a much bigger impact on a random read-and-write-intensive application, versus one that does mostly sequential reads and writes. Proper cache sizing, I/O queue depth settings, and cache partitioning can be equally important as drive type and configuration in tiered storage designs.

Always use the right software tools that can interface to and are supported by the platform’s firmware, and that provide enhanced levels of functionality to optimize and tune all drive and RAID combinations.

#8: Design the Dynamic Provisioning Storage Pool

To gain optimum performance and benefits using thin provisioning:

- ✔ Design the thin provisioning configuration with multiple I/O and/or purpose-specific pools.
- ✔ Understand that multiple pools may conflict with the goal of widest possible striping and simplified storage provisioning (via less pools to manage).
- ✔ Validate application requirements carefully through baselining. Think of each pool as a tier and use data classification for proper assignment.

#9: Create a Dynamic Provisioning Pool

When creating the thin or dynamic provisioning storage pool:

- ✓ Do not intermix the hard drives with different characteristics.
- ✓ Create the pool with a minimum of four RAID groups, with one LUN in each consuming all of the RAID group's usable space. Each RAID group should be behind a different back-end director pair.
- ✓ For dynamically adding more storage to the thin provisioning storage pool, use normal I/O capacity planning to determine the size of expansion.

#10: Plan and Design a Virtualization Project

Successful implementation of a virtualized storage environment and advanced techniques (such as a tiered storage strategy and thin provisioning) will provide significant business advantages and an improved return on investment (ROI).



Always use the professional services offered by respective vendors, because they know their technology, hardware and software well, and they can design, plan, and implement the project successfully. They should also teach you how to manage the project after the implementation is completed.

Through proper implementation of storage virtualization, IT can be more accountable and will provide business-level reporting on a common set of metrics. Many companies do not measure performance in a consistent way, especially with newer virtualization technologies. Few have good cross-domain coordination or have established an agreed-upon set of metrics for monitoring and managing performance expectations. Service level agreements (SLAs) can't be implemented without a good understanding of the metrics.

If the right metrics are put in place, with tools and processes to help measure and manage them, then performance SLAs can be established relatively easily.

