

Making Everything Easier!™

Hitachi Data Systems Edition

Storage Economics

FOR
DUMMIES®

Learn to:

- **Make a business case for storage technology investments**
- **Understand the four principles of Storage Economics**
- **Reduce storage costs using economically superior solutions**

Brought to you by

HITACHI
Inspire the Next

©Hitachi Data Systems

Justin Augat

Foreword by Hu Yoshida



About Hitachi Data Systems

Data drives our world — and information is the new currency. Send it. Receive it. Store it. Manage it. Process it. Protect it. Use it. Information has become our most valuable asset. The challenge for global businesses lies in making information both available and secure. Hitachi Data Systems solutions help organizations transform raw data into valuable information by delivering on a vision that IT must be: Virtualized, Automated, Cloud-ready, and Sustainable.

We don't sell just storage. We provide customers with solutions to make business better on every level. Our solutions manage data growth, and collect and connect data to create valuable information. They help cut costs and reduce operational complexity. They empower people to access, share, and secure information. They help organizations innovate, collaborate, and grow. They have the power to enhance and change lives.

Hitachi Data Systems is trusted by the global business community. Our customers include:

- 44 of the top 50 Fortune Global 500® companies
- 81 of the top 100 Fortune 1000® companies
- The top 16 global property and casualty insurers
- The top 15 global commercial banks
- The top 10 global telecommunication companies
- The top 10 global aerospace and defense companies
- The top 7 global technology companies

Storage Economics

FOR
DUMMIES®

HITACHI DATA SYSTEMS EDITION

by Justin Augat

Global Product Marketing, Hitachi Data Systems

Foreword by Hu Yoshida

Chief Technology Officer, Hitachi Data Systems



WILEY

John Wiley & Sons, Inc.

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

Storage Economics For Dummies; Hitachi Data Systems Edition

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

Copyright © 2012 by John Wiley & Sons, Inc., Hoboken, New Jersey

Published by John Wiley & Sons, Inc., Hoboken, New Jersey

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests 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 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. Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

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

For general information on our other products and services, please contact our Business Development Department in the U.S. at 317-572-3205. For details on how to create a custom For Dummies book for your business or organization, contact info@dummies.biz. For information about licensing the For Dummies brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN 978-1-118-17787-7 (pbk); ISBN 978-1-118-17797-6 (ebk)

Manufactured in the United States of America

10987654321



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

Table of Contents

Foreword	vi
Introduction	1
About This Book	2
How This Book Is Organized	2
Icons Used in This Book.....	3
Where to Go from Here	4
Chapter 1: Why Storage Economics Matters	5
Mapping the Global IT Economic Landscape.....	5
Money: The Language of Business.....	7
Some Simple Economic Truths	9
Chapter 2: What Is the Cost of That Computer in the Window?	11
The Price Quandary — Why Price Is Not the Same as Cost ...	11
Cost Comes in All Flavors	14
Costs have families, too.....	18
Correlating costs to different segments of the IT Department	20
Chapter 3: Know Your Costs	23
Identifying Your Costs.....	23
Establishing a Baseline.....	27
You Can't Improve What You Can't (Or Don't) Measure....	28
Looking Ahead.....	32
TKO-ing a TCO.....	32
Identifying the costs.....	32
Measuring the costs	33
Reducing the costs	33
Chapter 4: Implementing Economically Superior Solutions	35
Recognizing Economically Superior Solutions.....	36
Virtualization.....	36
Thin provisioning	36
Dynamically tiered storage.....	37
Active archive	37

- De-duplication..... 37
- Solid state drives 38
- Zero page reclaim 38
- Storage and storage area network consolidation..... 38
- Information Technology Infrastructure Library
 - best practices..... 38
- Disaster recovery..... 39
- Virtual tape library 39
- Disk-based backup..... 39
- Unified backup 39
- Unified console 39
- Policy-based management 40
- Storage architecture..... 40
- Storage services catalog 40
- Utility services 40
- Outsourcing..... 41
- Managed services 41
- Massive array of idle disks, spin-down disk..... 41
- Public, private, and hybrid cloud 41
- Mapping Costs to Solutions 42
- Formulating an Improvement Plan 45
- Checking Your Progress..... 46

- Chapter 5: Examining Some Economic Elixirs 47**
 - Virtually Everything..... 48
 - Improving Capacity Utilization and Performance with
 - Thin Provisioning..... 51
 - Aligning Data with Its Business Value 53
 - Benefits of the “Big Three...”..... 56

- Chapter 6: Show Me the Money 61**
 - Reducing Costs Isn’t the Whole Picture..... 61
 - The Storage Economics Hero 63
 - First, think like an economist 63
 - Second, talk like an accountant 63
 - Finally, act like a technologist..... 64

- Chapter 7: Ten Ideas for Reducing Total Cost
of Storage Ownership 65**

Publisher's Acknowledgments

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

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

Acquisitions, Editorial, and Vertical Websites

Development Editor: Elaine Marmel
Senior Project Editor: Zoë Wykes
Editorial Manager: Rev Mengle
Acquisitions Editor: Kyle Looper
Business Development Representative:
Karen Hattan
Custom Publishing Project Specialist:
Michael Sullivan

Composition Services

Senior Project Coordinator: Kristie Rees
Layout and Graphics: Claudia Bell,
Cheryl Grubbs, Lavonne Roberts,
Julie Trippetti
Proofreaders: Lauren Mandelbaum,
Dwight Ramsey

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

Kathleen Nebenhaus, Vice President and Executive Publisher

Composition Services

Debbie Stailey, Director of Composition Services

Business Development

Lisa Coleman, Director, New Market and Brand Development

Foreword

The world of Information Technology (IT) is entering a major new phase where information technology is everywhere, both inside and outside the data center. Until recently, IT drove business, but today, business is driving IT — and IT is falling behind. Businesses cannot wait a month to provision an application to exploit a new market opportunity, or a day to spin up additional resources to meet sudden peaks, or hours to analyze a shift in buying patterns. New technologies like virtualization and new ways to implement technology like cloud are being used to increase IT agility and scalability, with more new technologies appearing every day. We undoubtedly will be making investments in new IT technologies, but how will we be able to make the right investments and not be driven by the latest hype?

Storage Economics is a methodology developed by David Merrill, Chief Economist at Hitachi Data Systems, for identifying the total cost of ownership (TCO) within the storage infrastructure, which elevates the discussion beyond the cost of acquisition and looks at three- or five-year operational and capitalization costs. Storage Economics incorporates return on investment (ROI) and return on assets (ROA) analyses and enables IT managers to communicate with the financial decision-makers in terms that they understand. Finally, Storage Economics incorporates a project plan that identifies benchmarks in monetary units and includes measurements against those benchmarks.

This book introduces Storage Economics to both IT practitioners and to financial people and will help both types of professionals make better, measureable decisions that align IT with business.

Hu Yoshida
Chief Technology Officer
Hitachi Data Systems

Introduction

The lifeblood of almost every organization is data, and the requirements to protect, maintain, and access the systems that store that data are absolutely essential to support business operations. With the sheer amount and varying types of information that must be stored, the complexity of managing the data can grow as quickly as the amount of data itself.

In this fast-moving day and age, businesses cannot afford to wait for anything — if they do, they risk, at best, losing sales to the competition, or, at worst, falling behind the competition and being pushed out of business. The world of information technology (IT) is facing new challenges: IT planners are asking themselves, “How can we support the business model and simultaneously remain cost-effective?” Financial people are asking, “How can we identify the best and most cost-effective technology investments to make?”

Storage Economics, a methodology developed by David Merrill, Chief Economist at Hitachi Data Systems, helps answer the question of cost as it relates to the storage infrastructure. However, you can apply the principles of Storage Economics broadly across the entire data center and to IT in general. Using these concepts, IT professionals can learn to present business cases that justify investments in information technology and reduce existing costs with best practices. In addition to calculating the metrics typically used by financial professionals, such as return on investment (ROI) and return on assets (ROA), Storage Economics helps you determine the total cost of owning storage (TCO). Before you start a plan to reduce costs, you need to first be able to identify and measure all the costs that will occur over the life of the investment. Storage Economics makes the case that the cost of acquiring assets is just a small percentage of their total cost to own, and to accurately measure TCO, you must identify the operating and capital expenses associated with a storage purchase over a three- to five-year window.

About This Book

Since 1999, the folks at Hitachi Data Systems (HDS) have been defining and refining principles, concepts, cost categories, and models around the Storage Economics methodology to identify, measure, and reduce the costs of owning storage. This book introduces the concepts of Storage Economics from Hitachi Data Systems and should be valuable to two groups of people in helping them make decisions that align investments in information technology with corporate goals:

- ✔ IT practitioners who want to learn how to justify funding requests in terms with which financial people and upper management will relate
- ✔ Upper management and financial people who want to understand the most effective way to calculate the true cost of purchasing storage — beyond just its price

In the past few years, HDS staff has found that these principles apply not only to storage in its raw form, but also to adjacent infrastructure areas. Why? Because investments in data centers, including servers, networks, and newer technologies such as hypervisors, virtual desktop infrastructure (VDI), and converged platforms, all need to be approached with the same goal with which one approaches an investment in storage — to ensure that the right investments can be shown to reduce unit costs.

How This Book Is Organized

This book is organized into seven chapters as follows.

Chapter 1: Why Storage Economics Matters. This chapter delves into the financial conditions facing information technology departments and the basic principles of Storage Economics that help IT departments succeed in deploying a cost-efficient storage infrastructure.

Chapter 2: What Is the Cost of That Computer in the Window? This chapter looks at the 34 costs associated with storage that HDS has identified over several years of research, and groups these costs into families for easier identification.

In addition, this chapter outlines how the 34 costs of Storage Economics can be applied and mapped to various portions of the data center.

Chapter 3: Know Your Costs. Because each organization is unique, only some costs apply to any given business infrastructure. Use the information in this chapter to help you identify the costs your organization wants to reduce and learn about the best ways to measure those costs.

Chapter 4: Implementing Economically Superior Solutions. All solutions to reduce storage ownership costs are not created equally; some solutions are economically superior. Chapter 4 presents 22 examples of economically superior ingredients and describes the best approach to implementing them.

Chapter 5: Examining Some Economic Elixirs. This chapter looks at three economically proven solutions — the ones that provide the greatest reduction in storage TCO — in detail.

Chapter 6: Show Me the Money. You can do more than just reduce TCO with Storage Economics. This chapter explores some potential revenue-generating opportunities for IT departments. It also examines the real-world examples of people who have succeeded using Storage Economics from HDS.

Chapter 7: Ten Ideas for Reducing Total Cost of Storage Ownership. Here, you discover ten actionable ideas for reducing costs.

Icons Used in This Book

For Dummies books highlight certain kinds of information with icons to attract attention. The following is an explanation of those icons:



The text accompanying this icon is good to keep in mind for future consideration.



This icon indicates details that may interest the technically savvy.



When you see this icon, some pieces of information are offered to clarify certain specific points.



When you see this icon, something could go wrong, so make sure that you read the paragraph. This icon warns you of common mistakes and ways to avoid them.

Where to Go from Here

The chapters in this book are best read in order — particularly Chapters 1 through 4. Here's why. You'll have a better sense of Storage Economics and what it can do for you. In Chapter 1, you find out why Storage Economics even matters. And, Chapters 3 and 4 build on each other, going into the kinds of costs you can look to reduce, identifying the costs important to you, and looking at economically superior solutions and how to implement them.

But hey, it's your book so feel free to do what works best for you. Jump around wherever you'd like or simply turn the page and just start reading.

Chapter 1

Why Storage Economics Matters

.....

In This Chapter

- ▶ Understanding economic conditions for information technology
 - ▶ Realizing how businesses measure performance
 - ▶ Introducing the four principles of Storage Economics
-

Global, national, local, and business economic factors require that information technology (IT) departments constantly evaluate and implement plans to reduce costs. Even as the size of the IT department continues to grow in terms of storage capacity, servers, applications, and the amount of data being generated, businesses demand that IT departments hold costs flat and do more with fewer resources.

This chapter explores the financial conditions facing IT and the four principles of Storage Economics for reducing total cost of ownership (TCO) in the storage infrastructure.

Mapping the Global IT Economic Landscape

Over the last 60 years, the explosion of available data has driven the growth and importance of the IT department within the business organization framework. During the first wave of data growth, data management was centralized around computer mainframes. The second wave saw the arrival of the personal computer and a drive toward decentralization

and individual control over data. The third wave revived the need to share data but took advantage of new technology to network individual computers so that they could share data more intelligently. The current wave has seen the addition of new intelligent devices to the computing mix — devices like portable computers and cellphones — with new devices appearing all the time.

Even though data centers tend to double their storage needs every 18 to 24 months by some estimates, IT departments can no longer fill their needs *carte blanche*; every purchase is subject to review. Even though users need to store more data — increasing the demand for IT resources at the rate of 30 percent per year, again by some estimates — both capital and operational budgets are being cut or remaining flat (see Figure 1-1).

Lower acquisition costs are no longer sufficient; it's now essential to search for ways to reduce both capital expenditures (CAPEX) and operating expenses (OPEX). IT professionals need to make business cases to gain management support for both strategic and tactical investments.

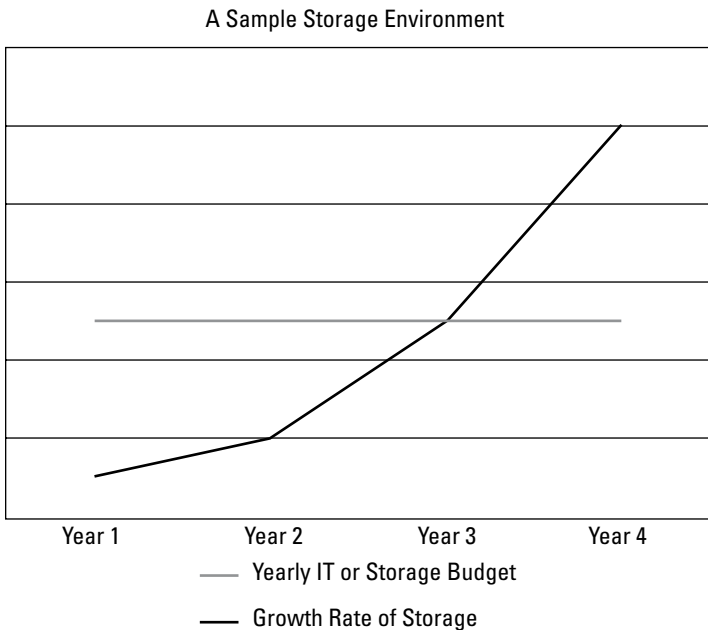


Figure 1-1: Comparing IT budgets to the demand for IT resources.

Money: The Language of Business

System engineers and information technologists have for many years believed that technology is supreme, a driving force within the organization, and above budgetary constraint. After all, systems must be purchased and upgraded, and the newest technology is essential to keep up with global demands and to maintain a competitive edge.

Ideally, the increased demand for IT resources (refer to Figure 1-1) — and the subsequent increased costs — would be the result of sales success, new revenue streams, or increased market share. Unfortunately, however, increased costs are more likely due to:

- ✔ Organizational change
- ✔ Increased competition
- ✔ Regulation, compliance, and data management
- ✔ Power, cooling, and (limited) floor space costs
- ✔ Increased cost of managing data
- ✔ Inefficient (or old) storage architectures

Growing the IT department depends on being able to sell a case for a particular investment to the people who hold the purse strings. Over the years, IT planners have had a harder time procuring funds due to increased budget scrutiny.

Many storage engineers and architects describe only the technological or operational benefits of a proposed initiative. But, adding a business focus can help you to sell your case. The financial decision-makers typically focus on three questions:

- ✔ How much is the investment?
- ✔ How quickly will we see a payback?
- ✔ How much is the net saving on the investment?

Further, the financial folks always balance investments in technology or operational efficiency against their ROI (return on investments), their potential to enable new business,

or their ability to establish a clear technical advantage. So, include in IT investment proposals a basic cash flow analysis and measure the ROI, ROA (return on assets), and TCO. When you recommend cost-reduction measures, identify clear metrics and provide a full analysis of the payback and savings of these initiatives.



Including a business focus in your analyses helps you to justify, in terms those outside the IT department can understand, the need to undertake the initiatives you suggest. If you can prove a business win in financial terms, you're more likely to get what you need than if you rely on technical jargon.

IT specialists can present measurable justifications to sustain and improve their computer domains by learning the language of business. Understanding and exploiting the marriage of IT concepts with financial or monetary concepts can be powerful leverage to get the right things done at the right cost. Storage Economics provides a framework to create such a marriage.

Hope is not a strategy (even if it is audacious)

Information technology has enjoyed an amazing run of becoming cheaper each year. Gordon Moore, co-founder of Intel, describes a long-term trend of computing hardware in a 1965 paper; his observation has become known as Moore's Law, and it states that the number of transistors that can be placed inexpensively on an integrated circuit doubles approximately every two years. Transistors weren't the only technology to demonstrate exponential improvements of various natures.

Moore's Law has been extended in the storage industry concerning the costs of hard disk storage per unit of information. While we thank Mr. Moore for his price reduction law, hoping for price erosion alone to solve your budget problems is a dangerous strategy. Prices of information technology are dropping, but the costs can run a contrary course. Cost-reduction tactics need to be tangible and actionable — hope is not a strategy.

Some Simple Economic Truths

Storage Economics aligns the operational and technical dimension of the storage infrastructure to a corresponding financial viewpoint. Storage Economics helps answer the basic question, “What is the best use of the storage investment dollar?” Many medium-to-large IT centers have used Storage Economics to achieve an economically superior way of deploying and growing their storage while simultaneously achieving measurable cost savings:

- ✓ Reduced waste
- ✓ Reduced labor time per terabyte of capacity
- ✓ Lower costs of growth
- ✓ Significantly reduced time and effort in migration, data protection, and disaster services
- ✓ Reduction in power, cooling, floor space, and maintenance fees for a unit of storage capacity

Storage Economics uses four key principles that can lead to continuous improvement and operational excellence:

- 1. Price does not equal cost.** The cost of storage includes far more than the acquisition price; in fact, the acquisition price accounts for approximately 15-20 percent of the actual cost to own storage. Hitachi Data Systems (HDS) has identified 34 types of costs that compose the total cost of storage ownership.
- 2. Storage ownership includes 34 associated costs.** Not all 34 costs are relevant to all organizations, so each organization must prioritize the costs to identify the ones that are most important to reduce.
- 3. There are economically superior storage architectures.** Some storage architectures are known to help reduce costs. Although some solutions have higher price tags, they can be cheaper to own since they can be mapped to the cost sensitivities of the organization as outlined earlier.

4. You can't improve what you can't measure.

Organizations should measure current costs and track progress in reducing them. Plan to build a deployment road map. You don't have to make all your changes at one time; you can deploy over time. And plan to measure cost savings at regular intervals to ensure that your plan is producing the results you expected.

Storage Economics delivers the following cost-efficiency benefits:

- ✔ Prioritizes cost types most relevant to the business
- ✔ Exposes total costs to managers and technologists equally, to assist in strategic plans and tactical investments over time
- ✔ Employs business case and payback methodologies to help financially and technically justify investments
- ✔ Recognizes key elements like people, processes, and technologies that your organization can apply to implement economically superior architectures
- ✔ Improves ROA by using existing systems more effectively — and therefore lowering IT spending
- ✔ Emphasizes better productivity and environmental efficiencies, which reduces the need for physical resources and helps to lower operating expenses
- ✔ Helps reduce both capital expenditure and operating expenses, which lowers TCO for storage assets
- ✔ Improves business agility to better manage growth and support change

These principles go beyond vendors, products, and data types. While technologies will certainly change over time, you can use Storage Economic principles to successfully control the costs of storage, to evaluate new technologies, and to guide planning and purchasing.

The rest of this book explores these principles in action.

Chapter 2

What Is the Cost of That Computer in the Window?

.....

In This Chapter

- ▶ Addressing the difference between price and cost
 - ▶ Examining the different kinds of costs
 - ▶ Grouping costs into families
 - ▶ Correlating costs to different segments of the IT department
-

If you expect to reduce storage costs, you need to know what your storage costs are. But before examining your costs, which takes place in Chapter 3, this chapter takes a look at the possible sources of storage costs that Hitachi Data Systems (HDS) has spent several years identifying.

All costs are not necessarily relevant to all organizations; that is, reducing one particular cost might have a major impact at one company and a negligible impact at another. So, the costs are organized into families since this approach often helps organizations focus on the costs that are important to them. And, these costs do not apply to just the storage infrastructure — they can be applied and mapped to several other segments of the IT department

The Price Quandary — Why Price Is Not the Same as Cost

The first of the four cost-reducing principles presented in Chapter 1 states that price is not the same as cost. Yes, there are some pretty big price tags for that new sales management

system or for expanding the network to hundreds of remote offices, but those price tags tend to account for approximately 15-20 percent of the long-term costs of owning those investments. Price erosion has played a big part in reducing the percentage that acquisition plays in the cost of an investment in storage.

Just a few years ago, the price of the hardware and software was a much higher portion of the total cost of ownership (TCO), but price erosion has changed that. Price erosion has been due, in part, to several factors:

- ✔ Consolidation of vendors
- ✔ Vendor competition
- ✔ Disk density
- ✔ Economies of scale from factory tooling and disk quantities

Over the past 40 years, storage designers have been shrinking storage media's grain diameters and decreasing storage media's thickness, putting Moore's Law in effect for storage media: Media capable of storing more data has become available for exponentially lower prices. This price erosion has driven down the total cost of acquisition (the procurement price) from 50 to 60 percent of the TCO to much closer to only 20 percent. To truly reduce the TCO, you can't rely on Moore's Law alone to provide lower price per terabyte.

Many storage decision-makers are seduced into thinking that using low-priced disk accomplishes their objective to lower the cost of storage ownership. In fact, purchasing inexpensive storage solutions most often leads to a higher TCO over a three- to five-year horizon. Low-priced disk options help with short-term capital expenditures (CAPEX), but will most likely lead to higher operating costs such as maintenance time and effort, and more frequent capital costs in the longer term.

What makes up the TCO? In many cases, labor may be the single largest cost element, followed by electricity, backup, migration and then the cost of acquisition. Migration costs, power/space/cooling, and all of the maintenance costs of hardware and software are the next three most dominant cost elements (see Figure 2-1).




Figure 2-1: A look at a simple total cost of ownership pie.

You can reduce TCO if you scrutinize the other 80 percent (nonpurchase part) of storage TCO — the costs other than acquisition. In many parts of the world, the cost of electricity is sufficiently high that the cost to electrify a storage array over three to four years will be more than the purchase price of the storage array. Labor costs are a substantial part of the IT budget, along with monitoring, protection, and decommissioning. Because the real costs of storage infrastructure are operating expenses (OPEX), managers need to identify these recurring costs and take action to reduce them. When preparing strategic plans, look beyond price to set in place cost-reduction investments. TCO is a more relevant measure as the price of storage approaches zero.

To achieve real economic value, IT planners have to look at the total cost of owning any particular IT investment. To calculate the TCO, you need to identify, measure, and then act on reducing the costs associated with IT investments. Some analysts predict that the cost of asset retirements has already surpassed the total cost of acquisition. Storage planners need to use TCO measurements to better isolate and measure cost efficiencies in the storage infrastructure.

Cost Comes in All Flavors

The process for reducing TCO involves

- ✓ Identifying costs to reduce
- ✓ Mapping the costs to solutions that reduce them
- ✓ Building a solution deployment road map
- ✓ Continuing to measure costs to ensure that they are going down

While it might seem like a daunting task, you *can* characterize the costs of storage. Table 2-1 lists 34 types of costs associated with the total cost of storage ownership. Some of these costs are hard or direct costs, others are soft or indirect. Some cost areas are OPEX while others are CAPEX.



Don't be overwhelmed by the idea of 34 types of costs for storage; you won't use them all. Why not? Because you don't want to boil the ocean. After reviewing the list, you'll find that many of the costs aren't relevant or important to your organization's infrastructure and you'll work with just the ones that will help you reduce your TCO.

<i>Cost</i>	<i>Description</i>
1. Hardware depreciation (lease)	Yearly costs for hardware depreciation or monthly leases
2. Software purchase or depreciation	Monthly or yearly costs for the purchase of the software. Some software can be capitalized with the original hardware acquisition.
3. Hardware maintenance	Recurring maintenance or warranty costs for all storage hardware after the base warranty period
4. Software maintenance	Recurring maintenance or warranty costs for all storage software
5. Storage management labor	Management labor costs associated with the various tasks of storage management, such as provisioning, tuning, load balancing, troubleshooting and upgrades

Cost	Description
6. Backup and disaster recovery labor	Aside from storage management, additional labor related to backups and restores, as well as disaster recovery planning and testing
7. Migration, remastering	Various costs associated with data migration at the end of the storage system's life. In large environments, there is continuous labor effort associated with data migrations. Remastering costs are associated with data life cycle costs while migration deals with the life cycle costs of the storage system.
8. Data mobility	Time and effort required to move data to different tiers or archive solutions. Different from remastering, data mobility follows the data life cycle, not the system life cycle.
9. Power consumption and cooling	kVA, BTU costs (converted to kW) of data center power. Power costs should include industrial-grade conditioning and battery or diesel backup.
10. Monitoring	Includes the cost of monitoring the storage area network, data network, console operations, and so on, using common protocols such as Simple Network Management Protocol (SNMP), Network Operations Center (NOC), and storage operations consoles.
11. Data center floor space	Cost per square meter of data center floor space. This often includes uninterruptible power supply (UPS) and raised floor costs.
12. Provisioning time	Business impact for the time waiting from when the request is made until capacity is presented to the host.
13. Cost of waste	Two types: usable but not allocated, and allocated but not used.
14. Cost of copies	Database management systems (DBMS) and other applications often require copies. In-tier or out-of-tier copies are possible. Test, development, quality assurance (QA), data mart, data loaders and similar applications all require multiple copies of structured and unstructured data.

(continued)

Table 2-1 (continued)

Cost	Description
15. Cost of duplicate data	Besides multiple copies, overhead is often associated with several copies of the same data. This is very common in unstructured file systems.
16. Cost of growth	Fundamentally, every storage architecture has a cost of growth. In high-growth environments with the wrong architecture, the cost of growth can be acute.
17. Cost of scheduled outage	Microcode changes, capacity upgrades
18. Cost of unscheduled outage (machine related)	In both the storage system and the connections or data path
19. Cost of unscheduled outage (people and process related)	Often due to capacity problems, operational control and physical thresholds
20. Cost of disaster risk, business resumption	Business impact with slow or fast recovery after a catastrophic event (declared disaster)
21. Recovery time objective (RTO) and recovery point objective (RPO) costs	Business impact costs resulting from the time it takes to return to a recovery time or point after a system failure or backup recovery
22. Data loss	Business and enterprise costs for lost, corrupted or unrecoverable data
23. Litigation, discovery risk	Legal risk and e-discovery time costs associated with lawsuits. This also covers general data location and recovery effort time.
24. Reduction of hazardous waste	Primarily a European Union cost due to regulations such as Reduction of Hazardous Substances (RoHS). Noncompliant hardware may incur an additional tariff for disposal of the asset.
25. Cost of performance	Impact to the business (good or bad) relative to total storage performance measured in input/output operations per second (IOPS), latency, and MB/second

<i>Cost</i>	<i>Description</i>
26. Backup infrastructure	Fixed cost infrastructure for backup. This includes backup servers, media servers, tape libraries, drives, etc.
27. Backup media	Local and remote media costs for backup; recurring and capacity related costs
28. Cost of risk with backup windows	Business impact of shortened or limited backup windows
29. Common Internet File System (CIFS)- or Network File System (NFS)-related infrastructure	Filers, gateways and the necessary software to provide file servers and shared services in the enterprise
30. Local and remote data circuits	Dark fibre used for storage area network extensions, remote replication, and the associated software
31. Storage area networking	The costs of the hardware and software components of the storage environment that communicate with and provide access to the physical data storage – including Fibre Channel, Internet Small Computer System Interface (iSCSI) or network attached storage (NAS) connection infrastructures. This cost also includes routers, gateways, host bus adapter switches, and directors.
32. Noncompliance risk (archive, data retention)	Data storage is subject to legal and legislative requirements, such as the Health Insurance Portability and Accountability Act (HIPAA) of 1996, Basel III, Sarbanes-Oxley, and carbon emissions legislation. Unless data are stored, protected, and preserved appropriately, organizations incur risks such as fines, negative publicity, and criminal prosecution.
33. Security, encryption	Costs associated with protecting, securing and encrypting data and the storage infrastructure
34. Cost of procurement	Costs associated with time and effort required to acquire hardware and software. This includes required preparation, review, negotiation, selection and certification processes.

A word on hard and soft costs

Hard costs are quantifiable costs that you can measure. But not all Storage Economics costs are quantitative; some, like the cost of risk or the cost of a future potential event, are soft costs that are qualitative in nature. Reducing soft costs can be part of an effective business case for change. When possible, converting technical and operational benefits

into monetary terms will help make the business case for a storage strategy. In fact, every storage action or reaction can be stated in economic terms. Soft costs that are hard to measure still need to be identified and measured in a structured way to present the full impact of all risks, benefits, and actions.

Costs have families, too

You can group the 34 types of storage costs into cost families; this approach often helps IT professionals focus on a particular type of cost. The 34 types of storage costs are organized into the following six major families. Please note that some costs are included in more than one family.

Procurement Costs

- ✓ Hardware depreciation (lease)
- ✓ Software purchase or depreciation
- ✓ Hardware maintenance
- ✓ Software maintenance
- ✓ Cost of waste
- ✓ Cost of duplicate data
- ✓ Cost of growth
- ✓ Cost of procurement
- ✓ CIFS-, NFS-related infrastructure
- ✓ Local and remote data circuits
- ✓ Storage area networking

Management, Operations, Life Cycle Costs

- ✓ Storage management labor
- ✓ Backup and disaster recovery labor
- ✓ Migration, re-mastering
- ✓ Data mobility
- ✓ Monitoring

Environmental Costs

- ✓ Power consumption and cooling
- ✓ Data center floor space
- ✓ Reduction of hazardous waste

Business Impact Costs

- ✓ Cost of scheduled outage
- ✓ Cost of unscheduled outage (machine related)
- ✓ Cost of unscheduled outage (people and process related)
- ✓ Cost of performance
- ✓ Provisioning time

General Risk

- ✓ Cost of risk with backup windows
- ✓ Litigation, discovery risk
- ✓ Data loss
- ✓ Cost of disaster risk, business resumption
- ✓ Recovery time objective (RTO) and recovery point objective (RPO) costs
- ✓ Noncompliance risk (archive, data retention)

Information and Data Protection Costs

- ✓ Cost of copies
- ✓ Backup infrastructure
- ✓ Backup media
- ✓ Security, encryption

Correlating costs to different segments of the IT Department

Finally, the 34 types of storage costs are not specific to a single component of the IT infrastructure. Typically, you find these costs in several different technology areas. Since the types of storage costs also differ by the infrastructure or functional IT element, Table 2-2 helps to show the correlation of company costs to different IT infrastructure elements.

	<i>Storage</i>	<i>Servers</i>	<i>Desktop, VDI</i>	<i>Network</i>	<i>Data Center</i>
1. Hardware depreciation (lease) expense	✓	✓	✓	✓	✓
2. Software purchase or depreciation	✓	✓	✓	✓	✓
3. Hardware maintenance	✓	✓	✓	✓	✓
4. Software maintenance	✓	✓	✓	✓	✓
5. General admin/management, labor	✓	✓	✓	✓	✓
6. Backup and DR labor	✓				✓
7. Migration, remastering	✓				
8. Data or workload mobility	✓	✓			
9. Power consumption and cooling	✓	✓	✓	✓	✓
10. Monitoring	✓	✓	✓	✓	✓

	<i>Storage</i>	<i>Servers</i>	<i>Desktop, VDI</i>	<i>Network</i>	<i>Data Center</i>
11. Data center floor space	✓	✓	✓	✓	✓
12. Provisioning time	✓	✓	✓	✓	
13. Cost of waste (or reserve capacity)	✓				
14. Cost of copies (or redundant assets)	✓			✓	
15. Cost of duplicate data	✓				
16. Cost of growth, expansion	✓	✓	✓	✓	✓
17. Cost of scheduled outage	✓	✓	✓	✓	
18. Cost of unscheduled outage (machine)	✓	✓	✓	✓	
19. Cost of unscheduled outage (people/process)	✓	✓	✓	✓	
20. Cost of disaster risk, business resumption	✓	✓		✓	✓
21. Recovery time (RTO) costs	✓	✓	✓	✓	✓
22. Data loss	✓			✓	

(continued)

Table 2-2 (continued)

	<i>Storage</i>	<i>Servers</i>	<i>Desktop, VDI</i>	<i>Network</i>	<i>Data Center</i>
23. Litigation, discovery risk	✓				
24. Reduction of hazardous waste — ROHS	✓	✓	✓	✓	✓
25. Cost of performance	✓	✓	✓	✓	
26. Backup infrastructure	✓				
27. Backup media	✓				
28. Cost of risk with backup windows	✓				
29. CIFS, NFS related infrastructure	✓	✓			
30. Local and remote data circuits	✓	✓	✓	✓	
31. Storage area networking	✓	✓		✓	
32. Noncompliance risk (recovery, data retention)	✓				
33. Security, encryption	✓	✓	✓	✓	✓
34. Cost of procurement	✓	✓	✓	✓	✓

Chapter 3

Know Your Costs

In This Chapter

- ▶ Selecting costs to monitor
 - ▶ Taking a baseline measurement of costs
 - ▶ Identifying cost metrics to monitor
 - ▶ Reducing your costs through technology
 - ▶ Getting TCO information from vendors to evaluate potential cost-savings solutions
-

Storage Economics helps you to address more than the price of storage; it focuses on the total cost of storage ownership. Chapter 2 describes 34 different kinds of costs associated with storage. Since each organization is unique, you won't necessarily be affected in great part by all 34 costs. Instead, using the 34 costs as a basis, you should identify the costs specific to *your* organization that you want to reduce.

This chapter focuses on helping you identify your costs as well as cost metrics you can use to take a baseline measurement of your costs. It also describes how you can get total cost of ownership (TCO) information from vendors to evaluate potential cost-savings solutions they want to sell you.

Identifying Your Costs

Chapter 2 explains why the price of storage and the cost of storage are two different numbers, and it presents the 34 costs of storage that Hitachi Data Systems (HDS) has identified over the last several years. Most IT and financial professionals do not use all 34 costs; instead, they select a subset of 8 to 10 costs that are important to them for any number of

reasons. For example, the selected costs may fall into any or all of the following areas:

- ✔ Are anticipated to be the ones that will reap the greatest immediate and long-term savings if reduced
- ✔ Costs that are related to actual budget and line-item reductions
- ✔ Reflect key management priorities like utilization or return on asset (ROA)
- ✔ Could be the foundation of a chargeback scheme



Choosing the types of costs to reduce is a personal decision.

After hundreds of evaluations, HDS determined that the majority of IT infrastructures have the potential to significantly reduce costs for every terabyte (TB) of usable disk capacity within the storage infrastructure. Implementing the right technology to reduce operating expenses (OPEX) can generate net savings that directly impact your company's bottom line.

Understanding the type of savings you can expect is important. Sure, new storage architectures will help you save money, but with a balanced investment in people, processes, and products, you can generate the following relative savings depending, of course, on what actions you take to reduce these costs:

- ✔ Waste reduction, approximately 20 percent savings
- ✔ Outage time reduction, approximately 10 percent savings
- ✔ Management labor effort, approximately 15 percent savings
- ✔ Migration cost savings of 20-40 percent
- ✔ Maintenance fees, approximately 10 percent savings
- ✔ Environmental costs, approximately 10 percent savings
- ✔ Miscellaneous operational efficiency, approximately 5 percent savings
- ✔ Other, approximately 10 percent savings

Discouraging wasteful behaviors

IT as an industry is at a tipping point with over-allocated resources, and the finance people in organizations express displeasure when you present them with ROA, consumption, or utilization metrics. Here are some mechanisms you can implement to discourage wasteful behaviors.

- ✔ Implement a chargeback system and, within the system, offer significantly different rates for efficient versus non-efficient resources. Yes, it's true that chargeback systems continue to be unpopular, even in the face of tough economic purchasing times, but a chargeback mechanism might work for your organization. See Chapter 6 for more information.
- ✔ Try using a show-back system, where monthly reports or management metrics reflect overall ROA for the IT infrastructure assigned to a group or department.
- ✔ Offer two different availability models: fast availability (hours, days) for efficient resources and slow availability (days, weeks, months) for non-efficient resources.
- ✔ Establish a capital review process that includes a department or line of business that reports on usage trends and use of deduplication methods to better improve what they already have.
- ✔ Move to a centralized allocation and storage pooling governance model to centrally redistribute unused capacity.
- ✔ Put in place proven technology to aggressively reduce waste (see Chapter 4 for information about economically superior solutions).
- ✔ Transform to a service utility model or subscription service to pay for capacities that are not consumed or allocated.

While the aforementioned costs represent typical sources of cost reduction, other areas are increasingly emerging for cost reduction from within the storage infrastructure:

- ✔ Hardware improvements can come from the latest hard drive technologies and from moving long-term storage from tape to disk. Replacing or refreshing old disk is often an easy method to reduce OPEX because:
 - Old disks are typically out of warranty and often require more monthly maintenance, resulting in higher costs than new disks.

- Power, space, and cooling requirements per TB are always much higher for older disk than for today's systems.
 - Some technical solutions may not be available for older disks.
 - Older disks can be less efficient relative to newer disks.
 - Older disks are more likely to be problematic or fail.
- ✔ Software used to manage storage is an increasingly important investment area, and you can reduce costs if you reduce the touch labor portion of storage ownership. Labor ranges from 10-40 percent of the cost of storage, so improving storage management, backup, tracking, assignment, and provisioning can reduce the current and future labor component of storage costs. Software management tools can help drive down the touch labor content and simultaneously reduce storage waste and outages.
- ✔ You can improve organization and people, usually the largest cost component, if you properly train and organize your employees and ensure that they work with standard processes according to directions established in your storage strategy.
- ✔ Investing in best practices and processes as defined by the Information Technology Infrastructure Library (ITIL) usually increases the effectiveness of managed storage and ensures that you apply the right efforts to the right data at the right time.
- ✔ New business models for storage can reduce storage costs if you make a behavioral and management commitment to the new business models.
- ✔ Regulatory compliance for efforts such as HIPAA and Sarbanes-Oxley require investment options that fall outside normal return on investment (ROI) considerations. You cannot overlook the cost of compliance in a storage strategy.



To identify your costs, use the costs presented in Chapter 2 to make some decisions about what costs are important to you and then select 8 to 15 of them for measurement.

Establishing a Baseline

Part of measuring costs is establishing a TCO baseline. You can determine the effectiveness of any cost-savings initiative in only one way: by comparing costs before and after implementing the strategy. The TCO baseline establishes your TCO for storage at a given point in time — the “before strategy implementation” cost.

But a TCO baseline can serve another purpose: It can help you identify the costs on which you should focus when choosing costs to reduce. As you measure your TCO baseline, you may become aware that some cost elements are actually much higher than you thought. Initially, you may not have focused on reducing those cost elements, but examining the TCO baseline could change your cost-savings strategy.

Figure 3-1 shows the kind of breakdown you can expect in a TCO baseline. Note that the price of storage is a relatively small part of the TCO. Power, cooling, administration, and transformation costs are usually at par with or higher than their associated acquisition costs. In general, you’ll find that acquisition costs make up approximately 20 percent of total costs. **Note:** These cost breakdowns can vary by industry, by geography, and by several other factors.




Figure 3-1: A TCO Baseline shows that purchase price is around 20 percent of total costs of ownership.

You Can't Improve What You Can't (Or Don't) Measure

Just to make sure that everyone is on the same page, you must measure the effectiveness of a cost-reduction initiative to ensure that you are making progress. You also use measurements to determine whether you are making better use of your assets; measurements also provide senior management with a reasonable business justification for new investments.

Many organizations focus on three primary metrics when measuring effectiveness, each of which is meaningful at the right time and in the right place (see Figure 3-2):

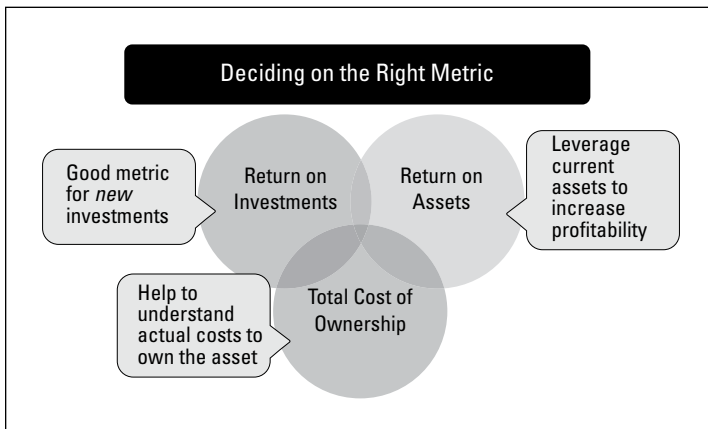


Figure 3-2: Deciding on the right metric.

- ✔ ROI is the metric you might choose for new investments: Essentially, do the cost savings justify the capital outlay? This measurement does not consider increasing the use of existing investments.
- ✔ ROA is a hard metric that can be difficult to measure. Why? Because ROA comprises several individual elements that, if improved, can contribute to an improved ROA metric. Performance, availability, environment, provisioning time, backup efficiency, and other elements all contribute to ROA. In what ways? Well, when you examine performance, you determine whether your hardware

is performing at its best levels. When you examine availability, you determine whether application downtime is minimized. Because ROA focuses on doing more with less and using your assets to their full potential, ROA must be the core message taken to the chief financial officer and to finance managers to justify investments to improve aggregate utilization.

- ✓ TCO measures the total cost to own your storage assets and compares them with alternative solutions and strategies.

In addition to these three metrics, you can choose from a variety of other metrics to help tell your storage cost progression story; some of the more popular metrics include

- ✓ For storage, TCO/TB/Year
- ✓ Total cost of data ownership (TCDO), dividing all the costs *not* by the usable capacity, but by the first instance data capacity
- ✓ TCO per transaction
- ✓ \$ per input/output operations per second (IOPS) per TB
- ✓ For greenhouse gas emissions, carbon emissions per TB
- ✓ For storage power consumption, TB per kilowatt hour (kWh) or TB per square meter of data center floor space
- ✓ Mean time to provision an application
- ✓ Recovery point objective (RPO)
- ✓ Recovery time objective (RTO)



The \$/TB metric is a popular sales technique used to compare vendors or track the downward trend of storage purchase price. Why? Because the metric doesn't require much information or analysis. But \$/TB is ineffective in terms of illustrating long-term ownership cost differences; the least expensive \$/TB solution could be very expensive to own if, for example, it repeatedly fails, requires large amounts of management time, and performs poorly for the applications.

Far too often, procurement is determined by the lowest storage bid, without considering that low-cost architectures often produce higher costs of ownership. Personnel responsible for (and measured on) acquiring lowest cost solutions do not

have to answer for the total costs incurred in the IT department. If, in the long-term, you want to reduce OPEX and CAPEX (capital expenditures), then you must use a total cost model. You can expect effective and continuous improvement when you measure and report on the right metrics.

Difficult-to-measure costs

Whenever possible, converting technical and operational benefits into monetary terms will help make the business case for a storage strategy. In fact, every storage action or reaction can be stated in economic terms. That said, not all Storage Economics actions are quantitative; some are qualitative in nature — we call these *soft costs*. Even though soft costs are hard to measure, you still need to identify them and measure them in a structured way to present the full impact of all risks, benefits, and actions. For example, reducing the cost of risk of an outage can be part of an effective business case for change — despite the fact that “risk” can be very difficult (if not impossible) to measure.

The fact is, the discipline of economics is largely based on non-measurable metrics. But not being measurable doesn’t make these metrics any less valuable. The valuation of a company, or more formally, the “market capitalization” of a corporation, is a good example of a valuable, nonmeasurable metric.

“Market capitalization” is calculated by multiplying the number of outstanding shares of a corporate stock by the share price. For example, 100,000,000

outstanding shares at a \$5 share price would equal a \$500,000,000 market capitalization. But, this calculation doesn’t really measure the valuation of a company.

If a coordinated event were held to sell every outstanding share of the corporation at the same time, or even a large percentage of the shares, the market capitalization would change due to the share price impact of such a large sell. More importantly, this represents just one way to determine the market value of a corporation. Other metrics exist, such as “enterprise value,” “book value,” “breakup value,” and more — and these calculations can vary by the method the investor or analyst chooses to use.

In any case, these metrics provide significant value to the user no matter how they are derived, and many people rely on them heavily to make important business decisions.

We can apply this approach to the measurement of storage costs. Are there storage and management metrics that are nonmeasurable? Absolutely. Are there macroeconomic metrics that don’t require any real counting? Yes. But are they any less valuable? Not necessarily. While many of these metrics can prove to

be subjective, often they are the best measurements we have. The trick is to understand why they can be subjective and why some measurements are better than others. Here are a few examples:

- ✔ **TB-per-person under management.** This metric seems to be simple; divide all the capacity between the number of people supposedly managing storage. But, this metric is full of holes since some of the storage administration team, if there is a dedicated team, also tends to look after servers, backup, data recovery tests, development, compliance requirements, and more. Fuzzy metrics like this one can be used at best internally to show how much more storage has been bought without adding headcount. The metric does not fairly represent the effective management nature of the staff and the data.
- ✔ **Total cost of storage.** Although TCO is the best metric you can use to calculate your costs, this metric is incomplete because it doesn't include the cost of the information on the disk. TCO of storage is the container cost, but if the container is largely empty or full of copies or aged data, then the metric is not helpful. We have to include the cost of the

first instance of information in the measurement; then TCO per TB is very measurable and very useful: The sum of all the costs divided by the total raw or usable capacity. The TCDO is more difficult to measure, but it is a far more meaningful measurement to the data owners.

- ✔ **Total cost of acquisition (TCA).** This measurement seems simple enough, since we can ask the procurement department to determine what was spent to buy a particular quantity of storage. But, once again, we would be measuring containers, and low cost containers may or may not adequately cover the business needs. Advanced containers cost more, but may do more (and ultimately cost less) in terms of operational, data protection, migration, and discovery functions for the business. In addition, the TCA is a small percentage — sometimes less than 20 percent — of the TCO and, as a metric, is fundamentally incomplete because it doesn't consider the maintenance costs of ownership over the life of the container. In fact, some analysts predict that the cost of retiring assets has already surpassed the TCA.

Looking Ahead...

Once you've isolated and measured your current costs, you can begin to identify *levers* — solutions that can drive costs up or down — that will help you reduce costs. Levers come in varieties, including but not limited to: technology, organization, procurement, business operations, and people skills.

Check out Chapter 4 for good cost-saving solutions you can consider, along with ways to approach implementation and measure progress.

TKO-ing a TCO

When you create requests for proposals (RFP) or requests for information (RFI) to which you want vendors to respond, ask for TCO data. You need vendors to give you more than price information; you need the total cost of any proposed solution. The best requests:

- ✓ Identify the costs that are important to you
- ✓ Provide information about your organization that the vendor can use to calculate TCO
- ✓ Require that vendors show how their solutions reduce your costs

Identifying the costs

As described in Chapter 2, Storage Economics includes 34 different types of storage-related costs. Since only some of those costs are relevant to each organization, you need to identify the cost elements on which you want to focus to effectively build a TCO. The most popular costs tend to be:

- ✓ Hardware depreciation
- ✓ Software depreciation
- ✓ Hardware maintenance after warranty
- ✓ Software maintenance after the warranty period
- ✓ Cost of growth or upgrades

- ✓ Power and cooling
- ✓ Floor space
- ✓ Administrative labor
- ✓ Migration costs
- ✓ Backup and data protection

Measuring the costs

Once you've identified the types of costs important to you, you need some local parametric data to create or measure the TCO. For example, if you include power and cooling, then you need to know the local cost per kilowatt hour of electricity. If you include administrative labor, then you need to know the annual, fully burdened — that is, including overhead costs — rate for labor.



As business types of people will tell you, the cost of labor isn't just the salary paid to an employee. Labor costs also include overhead amounts like benefits paid to the employee — think “health insurance” here — and payroll taxes paid for the employee.

For your potential vendors to respond effectively to your RFP or RFI, they need to know basic information about your location, capacities, growth, and asset inventories. They also need to know what costs are important to you so that they can recommend solutions that address your needs.

Reducing the costs

When responding to the request, vendors need to show how their solution will reduce what you have identified as costs that are important to you, so that the vendors can correlate their solution to those costs. You can't expect vendors to be clairvoyant and guess the types of costs you want to reduce in the TCO process, but if you supply that information, vendors should be able to map their solution to your costs.

Figure 3-3 shows you what a simple model of a six-year TCO for a proposed storage replacement might look like and includes the following types of costs: depreciation expense, maintenance

(both hardware and software), professional services, power and cooling, floor space, and labor for management.

Customer RFP Response							HITACHI DATA SYSTEMS
Six Year TCO Summary							
	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Total
Depreciation	163 200	163 200	163 200	163 200	163 200	-	816 000
HW Maintenance	0	0	0	76 296	76 296	76 296	228 888
SW Maintenance	0	0	0	33 828	33 828	33 828	101 484
Install Integrate	111 658						111 658
XX Power	43 289	43 289	43 289	43 289	43 289	43 289	259 734
YY Power	21 050	21 050	21 050	21 050	21 050	21 050	126 297
XX Space	2 160	2 160	2 160	2 160	2 160	2 160	12 960
YY Space	1 200	1 200	1 200	1 200	1 200	1 200	7 200
Administration	187 500	193 125	198 919	204 886	211 033	217 364	1 212 827
Yearly Cost	\$530 057	\$424 024	\$429 817	\$545 909	\$552 056	\$395 186	\$2 877 048
Yearly TCO/TB	\$1 409 73	\$1 127 72	\$1 143 13	\$1 451 89	\$1 468 23	\$1 051 03	\$1 275 29
						Blended 6 year average TCO/TB/Year	

Assumptions	5 years of depreciation	\$1 000 000 HW purchase price
	32% Marginal tax rate	\$200 000 SW purchase price
	6 years TCO model	
	\$125 000 per person labor costs, fully burdened	
	3% per year cost of labor increase	
	0% growth in storage capacity	
	\$0 15 kWatt hour for XX and is constant for 6 yrs	
	\$0 10 kWatt hour for YY and is constant for 6 yrs	
	\$60 00 square meter/month floorspace in XX	
	\$50 00 square meter/month floorspace in YY	

Figure 3-3: The kind of information you should expect from an RFP response.

When vendors respond in this manner, you are in a better position to compare vendor responses; you can measure the efficiency of any vendor’s design by the cost per TB/year. You can make side-by-side price and cost comparisons that will help you find the architecture that has the lowest cost of ownership for the price.



People often confuse TCA with TCO. There is nothing wrong with TCA — it tells an important CAPEX story — but TCA may be only 15-20 percent of the TCO. In the example shown in Figure 3-3, the labor will be more expensive than the hardware and software purchase price, and the cost of power/cooling will be about half of the purchase price.



Are you looking for the best recipe to purchase systems and solutions that provide critical functionality at the best price? Create an effective RFP or RFI that requires vendors to provide TCO information. Be sure to identify your costs, provide parametric data to measure the costs, and require the vendor to show you how their solutions will reduce your costs.

Chapter 4

Implementing Economically Superior Solutions

.....

In This Chapter

- ▶ Recognizing economically superior solutions
 - ▶ Mapping costs to cost-reduction strategies
 - ▶ Formulating a cost-reduction implementation plan
 - ▶ Monitoring costs over time
-

Businesses want to improve profits, and nobody disputes that reducing the total cost of ownership (TCO) for storage is one way to improve profits. Earlier chapters in this book outline the types of costs you can target for reduction and describe ways for you to take a baseline snapshot of your TCO before you make changes. This chapter tells you how to identify solutions that can help you reduce costs — and there are many such solutions. Since each organization's infrastructure is unique, all solutions don't work equally well. To identify the solutions best for your organization, you need to map the costs you want to reduce to solutions that will help you reduce them.

As you review solutions, you not only discover that some solutions actually help you to lower more than one cost element, but you also find that some solutions are cheaper to purchase. Remember, however, that the price of a solution doesn't equal the cost to own the solution. When you consider a solution, remember to determine whether it is economically superior to other possible solutions. Also, be aware that you don't need to implement all solutions simultaneously. You can plan to spread the implementation costs over time while you continue to monitor the cost-effectiveness of the solutions you choose.

Recognizing Economically Superior Solutions

Hitachi Data Systems (HDS) recognizes that you have many choices for increasing available storage capacity. HDS has also determined that some solutions not only increase available capacity but also help to reduce the TCO for storage and improve performance. HDS calls these solutions “economically superior.”

You can implement economically superior storage strategies either incrementally over time or as part of a complete data center technology refresh. The solutions you choose may not be the cheapest to buy, but they can be the cheapest to own. Some are best suited for high-growth environments, others reduce compliance or legal support costs, and others are best if power and floor space are your major concern; remember, each organization’s infrastructure is unique.

Virtualization

Storage virtualization — which comes in many types — is a powerful solution that consolidates, simplifies, and optimizes multiple storage devices in the data center so they can be viewed and managed as a single pool of storage, grow efficiently in line with business requirements, and reduce business downtime and application disruption. Array-based storage virtualization allows *heterogeneous arrays* (from many vendors) to be consolidated and to inherit the best features of the parent array. This technology can repurpose and extend the useful life of older storage arrays from different vendors and help defer future capacity purchases while significantly lowering capital and operational costs.

Thin provisioning

Traditional storage provisioning entails allocating on a “just in case” basis so that applications don’t run out of storage. Unfortunately, since applications don’t always use all the storage allocated to them, traditional provisioning can cause massive storage waste in the data center.

Thin provisioning pools storage and automatically and efficiently manages it for applications. This ensures that applications have enough storage when they need it — but not so much that it is wasted — effectively optimizing your available storage and significantly improving performance. This “just in time” instead of “just in case” provisioning allows administrators to reclaim allocated but unused storage so that you can improve storage utilization and reduce or defer future capacity purchases.

Dynamically tiered storage

Not all data is created equal. Some data can be classified as business critical, like sales and customer relationship management (CRM) data. Other data, such as human resources or payroll, may be classified as less critical. Storage media can also be classified according to business requirements. For example, solid state is fast, expensive, and appropriate for I/O intensive applications, whereas SATA (Serial Advanced Technology Attachment) is less expensive but slow and suitable only for non-critical applications.

Dynamically tiered storage automatically classifies and moves data to the right storage media based on its business need. The impact on administrators is two-fold: They can reduce the total cost of storage by using cheaper storage for non-critical and old or stale data, and they can dramatically improve application performance by ensuring that the appropriate storage media is used.

Active archive

All data has a life cycle associated with it. As data becomes inactive or accessed less frequently, you don't need to store it alongside active or critical data. An active archive is a dedicated storage system for the long-term preservation and protection of data as it nears the end of its life cycle. An active archive requires no backup and significantly less management, reducing the overall cost of storing, managing, and protecting stored data in line with its value to the business.

De-duplication

In most cases, mass amounts of duplicate data reside within the same data center. Copies and versions of the same files

persist throughout and have contributed to the data growth experienced by most companies. De-duplication technology identifies, manages, and reduces duplicate files to one version in a manner that is transparent to the author or user, reducing storage needs and improving utilization.

Solid state drives

As data is classified based on business value, some data needs to reside on the highest-performing storage media. Solid state drives (SSD) are designed with business-critical applications in mind. They provide high-performance input/output (I/O) and reduce access time, latency, and power usage.

Zero page reclaim

Even within an efficient storage architecture there is room for waste reduction. Zero page reclaim is a technology that identifies unused portions of storage capacity such as pages with zeroes resulting from deleted data and restores those unused portions back to the storage pool. Zero page reclaim improves the amount of storage that is actually used by applications — the storage utilization rate.

Storage and storage area network consolidation

Most data centers have multiple storage arrays and networks, resulting in a decentralized architecture that can be cumbersome to manage. Storage and storage area network (SAN) consolidation helps to centralize multiple arrays and networks and reduce the time, effort, and expense of managing them in aggregate.

Information Technology Infrastructure Library best practices

Information Technology Infrastructure Library (ITIL) is a widely adopted approach to IT Service Management that

provides a practical framework and best practices for identifying, planning, and supporting IT services to the business.

Disaster recovery

Part of the job of a storage administrator is to prepare for unforeseen shocks in the system that could disrupt business and impact revenue. *Disaster recovery* refers to the process, policies, and procedures used to facilitate preparation for recovery or continuation of the storage infrastructure in the event of a disaster or business disruption.

Virtual tape library

Virtual tape library (VTL) is a data storage technology used for backup and recovery purposes. VTL uses the latest disk and tape technologies to emulate legacy physical tape libraries, enabling staff to back up applications without having to rip and replace existing software.

Disk-based backup

Many data centers still use tape media for their backup and recovery needs. As the price of disk declines year over year, many data centers are implementing newer disk-based backup technologies to reduce the management and cost of backing up.

Unified backup

Many infrastructures consist of both physical and virtual backup strategies, and managing multiple systems can be time consuming and expensive. Unified backup is a technology solution that enables you to manage physical and virtual backups together through a common backup utility and infrastructure for virtual servers, mainframe, virtual desktop infrastructure (VDI), and so on.

Unified console

Multiple storage arrays and data types within the data center mean complex and time-consuming management, administration, and reporting. Unified console is a solution that simplifies

these tasks by presenting the storage environment from a logical and central view.

Policy-based management

As complexity grows alongside data, establishing an administrative approach to common situations can be very beneficial. Policy-based management is used to simplify data storage management by establishing policies and procedures to deal with situations that are likely to occur.

Storage architecture

There are many types of businesses — and many types of storage architectures that are best suited for those businesses. Some examples of storage architectures include SAN, NAS (network-attached storage), or unified architecture. Any storage architecture should be implemented based on business and data needs. The appropriate architecture standardizes solutions and provides a referential framework for high levels of repeatability, manageability, and serviceability.

Storage services catalog

To create an efficient and manageable storage environment, you want to align service and business objectives to the available storage tiers and capabilities, including remote access service features and quality of service — an often difficult task. You can take advantage of a storage services catalog to help you achieve these goals.

Utility services

Many data centers are reducing capital costs related to hardware and software purchases by employing utility services. A utility services storage model provides a “pay per use” or “pay for what you consume” framework that can be developed in-house for adjacent business units or purchased through a third-party provider.

Outsourcing

Outsourcing is a solution that offloads some part of the data storage environment, such as assets, labor, or applications, to third-party control. Many IT operations are choosing outsourcing as a cost-reduction strategy as capital costs come under scrutiny.

Managed services

“Managed services” is a hybrid solution of outsourcing and ownership. With managed services, data centers retain control and management of some portion of the infrastructure, such as mission-critical operations, and partner with third-party vendors to manage other parts, such as inactive data archiving.

Massive array of idle disks, spin-down disk

Massive array of idle disks (MAID) technology can power down drives or array groups using rules or policies based on access patterns. Less frequently accessed drives can sit idle more often. This technology can impact power- and cooling-sensitive environments such as those found in metropolitan areas and help to reduce operating costs.

Public, private, and hybrid cloud

Cloud computing has quickly become a widely available and talked about strategy for data center cost reduction. The public cloud has become prevalent in recent years to provide storage and/or applications to the general public over the Internet so that users can access data from nearly any location. The public cloud ensures reliability by providing, in some cases, significant data protection capabilities as part of the service cost.

A private or hybrid cloud is operated to provide specialized storage services where the cost or security risks are lower than in-house capabilities. Hybrid cloud solutions contain resource and infrastructure both behind and external to the company firewall.



Although you'll probably mix price and cost when evaluating these solutions, keep in mind that low-price storage architectures hit a point of diminishing returns and become more expensive to own than counterparts that were more expensive to buy but, over time, prove less expensive to own.

Mapping Costs to Solutions

Since each organization's infrastructure is unique, all solutions are not necessarily appropriate for all organizations. To identify the solutions that will give you the most bang for your buck, you need to map solutions to the cost areas you want to reduce — the ones you identify in Chapter 3.

You can use the TCO baseline information you collect to coordinate storage design, to prioritize investments and initiatives, and to create a road map of activities based on their projected potential to both reduce costs and support business needs.

The initiatives you identify may be strategic or tactical and will vary widely, depending on your environment. By mapping costs to cost-savings solutions, you can identify the multiple relationships between costs and potential cost-saving investments. Because one solution may serve to lower costs in more than one area, the mapping process helps you to identify the technology investments that should be your highest priority. Figure 4-1 presents a graphical view of this mapping.

You can map solutions to costs using this tool: www.hds.com/go/cost-efficiency/mapping-tool/. For example, you can map storage virtualization to the following costs:

- ✓ **Hardware/software maintenance.** Software maintenance costs often drop as virtualized assets inherit the controller's features and functions. Virtualization causes asset consolidation, enabling you to retire older arrays and save hardware maintenance costs.
- ✓ **Storage management labor.** The costs associated with heterogeneous storage management, such as provisioning, tuning, load balancing, troubleshooting, and upgrades often drop 5-10 percent.

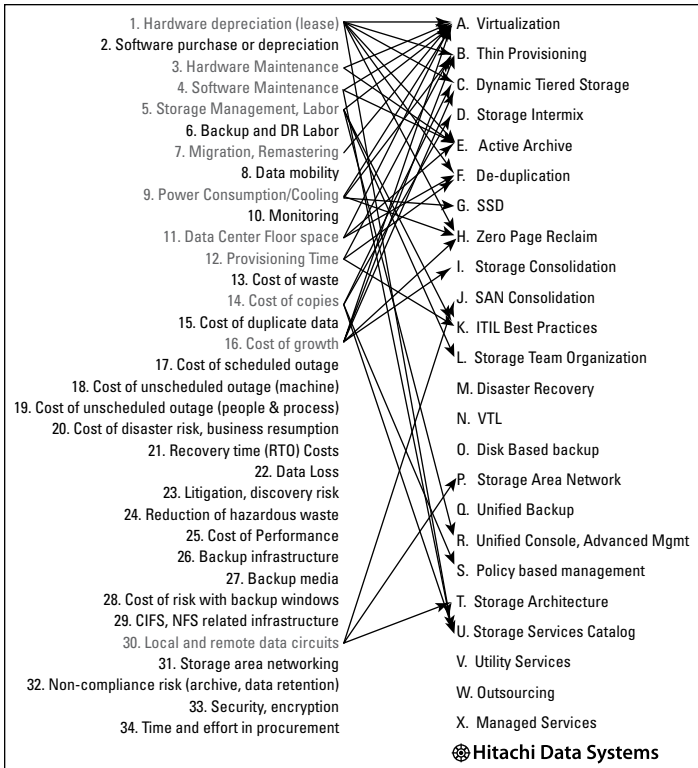


Figure 4-1: You need to map solutions to the right costs.

- ✔ **Power consumption and cooling.** Environmental costs go down because of better utilization, reduced waste, and the number of arrays you must support.
- ✔ **Data center floor space.** Asset consolidation and reclamation reduces these costs by better utilizing assets and reducing the number of arrays you must support.
- ✔ **Cost of waste.** Storage virtualization reduces two types of waste by 20-40 percent: usable but unallocated space and allocated but unused space.
- ✔ **Cost of migration.** By making hosts unaware of the movement of applications and applications unaware of the movement of data, virtualization can reduce migration costs by 70-90 percent.

- ✔ **Cost of growth.** The compounded impact of virtual storage, tiering, and thin provisioning can fundamentally reduce the appetite curve for storage capacity.

When mapping, avoid these pitfalls:

- ✔ Don't take a blind approach and select a single technology in the hope that it will reduce costs. Hope isn't a bad thing, but it also isn't a guarantee that costs will go down.
- ✔ Produce a plan that coordinates business needs and investments. For example, if the cost of power is not a significant element of your storage TCO, don't look for solutions that focus on reducing the cost of power.
- ✔ Don't just transfer costs with no real net reduction in unit costs. Simply shifting the costs without a fundamental re-architecting of a solution might not reduce total cost as you would expect.

That last pitfall bears a little further examination. You see, it's entirely possible to decide upon a cost-reduction initiative that does, indeed, reduce costs — just not *your* costs. Increasingly, medium-to-large IT shops are outsourcing IT operations to third-party contractors, system integrators, or cloud providers. Further, when it comes to storage, servers, virtual machines, and backup infrastructure, you can enter into a variety of agreements that affect asset ownership, labor allocation, task allocation, and more. So, when you embark on a cost-reduction plan, you need to determine whose costs your plan will target and impact — and who will save money in the end.

For example, if your system integrator or outsourcer owns and manages some of your storage capacity, a plan to increase storage utilization might reduce costs — but not necessarily for you. If you want to eventually see the savings, you may need to re-negotiate your contract with your third-party provider. Most organizations want to reap what they sow — and recognize the cost savings to their own budget before helping an external provider.



Make sure, before you implement a cost-saving initiative, that you determine exactly who will benefit by the initiative.



HDS uses a scorecard that aligns and correlates storage TCO with architecture options; the scorecard maps the costs you identify to potential storage investment categories. HDS then rates the impact of the investment on the costs and uses a weighted score to help you prioritize investments needed in a long-term road map to reduce storage costs.

Formulating an Improvement Plan

Undoubtedly, you'll identify more than one initiative you can undertake to reduce your costs. But, as is the way with much of life, you can't do everything at once. First, you probably won't have the money, and second, even if you had the money, it's rarely wise to change *everything* at the same time.

So, create a cost-reduction road map. Take a look at the impact of each cost-savings measure on the TCO baseline over the near and medium term. Then, you can build a road map like the one shown in Figure 4-2 that shows the impact of each solution on each of your key cost areas year by year.

A structured design and cost-planning approach like a storage improvement road map provides a key management tool and technical framework that you can use to prioritize investments.

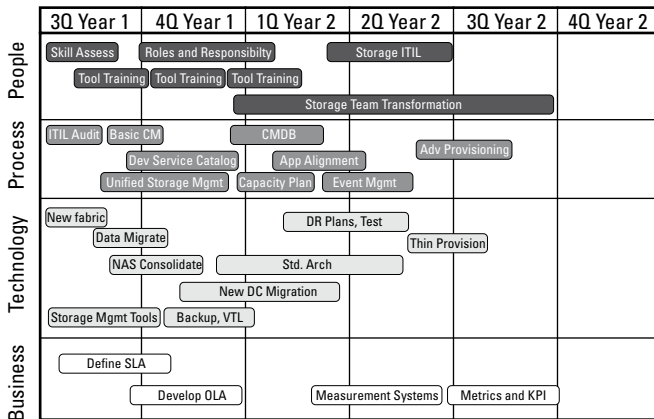


Figure 4-2: Establish a storage improvement road map.

Checking Your Progress

While this might seem to be stating the obvious, in this case, the obvious is worth stating: Once you start implementing the improvements you identify, you need to make sure that they are, indeed, lowering your costs. If you can prove to management that cost-reducing investments are truly reducing costs, you can use the information to help you justify further investment in IT.

Your TCO baseline comes into play for this exercise because it presents your TCO before you implemented any cost-savings strategies. You need to measure the same costs each year to determine whether your costs are coming down. Figure 4-3 shows the type of decline you might expect in TCO/TB/Year.




Figure 4-3: A sample TCO measurement over five years after implementing cost-saving measures.

Chapter 5

Examining Some Economic Elixirs

In This Chapter

- ▶ Taking a look at virtualization
 - ▶ Checking out the benefits of thin provisioning
 - ▶ Using automatically tiered storage
-

Economically superior storage architectures use technologies that maximize storage efficiency, reduce the capacity “appetite” curve, increase the value of your existing storage investments, and improve infrastructure performance assets. Among the 22 economically superior solutions described in Chapter 4, several stand out because they provide high performance, significant cost reduction, and scalability. At present, the “big three” elements of an economically superior storage architecture are storage virtualization, thin provisioning, and automatically tiered storage.

This chapter focuses on “the big three” and illustrates the individual and, when used together, compounded impact these solutions can have on total cost of ownership (TCO).



Just because a technology alternative (or a vendor of a particular technology alternative) promises to reduce your costs doesn't mean your costs will automatically go down. When evaluating a technology alternative, make sure that you identify its costs, calculate a TCO, and weigh the technology's costs against its benefits.

Virtually Everything

Virtualization has broad applications beyond the world of storage, but to reduce the TCO for storage, focus on storage virtualization. Storage virtualization is, in many cases, the foundation of economically superior storage architectures.

Imagine using one solution to manage all types of storage media and data, regardless of classification or file type. Then, imagine that nimble storage platform moving, without disruption, all your data across internal assets as well as externally attached assets from multiple vendors. There's no doubt that such an environment would improve functionality, utilization, and efficiency, and simultaneously reduce costs.



As fun as imagining is, it isn't a viable way to sell a cost-reduction solution to those who hold the purse strings. Although Hitachi Data Systems (HDS) has done TCO calculations, you can and should make your own calculation for your environment. You can't improve what you don't measure.

Storage virtualization uses software and supporting hardware to consolidate multiple storage systems in the data center and to simplify the inner workings of these systems to present a single, logical view. Essentially, virtualization enables administrators to decouple data center applications from the underlying storage, reducing complexity and creating a more efficient and performance-oriented environment.

You can consider several different types of storage virtualization, such as controller/array, file, network, and so on. In the case of controller/array-based storage virtualization, you can effectively aggregate multiple vendor storage systems seamlessly and manage them using a single interface. Controller/array-based virtualized storage also helps you to:

- **Significantly improve storage utilization.** Virtualization increases the aggregate amount of existing capacity available for provisioning by reducing or eliminating “data silos” (available storage capacity not accessible to the server/application that needs it).
- **Efficiently and non-disruptively migrate applications and data.** Because virtualization creates a layer between the application and storage, data can be moved seamlessly

from array to array and without disrupting application operation.

- ✔ **Extend the useful life of older assets and assets from multiple vendors.** With some types of virtualization, administrators can get more use from their storage arrays because they can perform technical refreshes quickly with little to no business disruption.
- ✔ **Improve and repurpose existing assets.** In addition to extending the useful life of assets, virtualized older assets can inherit the best capabilities of the parent array, enabling administrators to repurpose older, slower arrays to other functions that are appropriate for them instead of decommissioning them.

While there are many virtualization techniques, HDS virtualizes storage using storage controllers; this approach separates storage controllers from the storage media in any type of environment. No matter how you view virtualization or which form you choose, without a doubt it is the primary ingredient of “economic elixirs” that make an impact on TCO.

So what does that cloud look like to you?

Low asset utilization is forcing IT and business professionals to look at new options to improve return on assets (ROA), and cloud offerings are the hot topic of the day. The economics of the cloud play an important role in long-term funding, costs, and strategic planning. But, before you decide that clouds are a solution to reduce costs, make sure that you know what kind of costs you’re interested in reducing.

Virtualization is a foundation of cloud services, and many people use delivery options to describe how you will consume virtual cloud resources:

Note: Virtualization is a key enabler of cloud services, but virtualization is not exclusive to the cloud; you can use virtualization outside of the cloud infrastructure to reduce costs.

- ✔ Private clouds use economically superior solutions such as virtualization, archiving, advanced backup, thin provisioning, and de-duplication. But private clouds still require local effort and resources such as power, cooling, and management labor to operate the assets. From a cost perspective, many private cloud offerings differ very little

(continued)

(continued)

from current “capitalize-and-own” operational models.

- ✔ Hybrid clouds move some but not all costs to an external resource, blurring some of the operational cost lines that exist now with traditional IT services. When engineering the hybrid cloud, differentiate the costs you retain from the costs that you hand over to an external source. Why? Because the costs you hand over to an external source become additional costs that you haven’t incurred in the past. It’s entirely possible in hybrid cloud solutions that the total gross costs *increase* — instead of decrease. This increase may not be entirely bad news because your organization could gain a real business benefit from the added technical flexibility that clouds offer — a benefit that outweighs the added cost.
- ✔ Public clouds also use economically superior solutions such as virtualization, archiving, advanced backup, thin provisioning, and

de-duplication. In a public cloud, many organizations share these technologies and their associated costs. The economies of scale of a public cloud offer the lowest cost, but the public cloud comes with some new and often high risks, such as the general risk associated with moving data, processing, or applications to the external world; cloud outages highlight this risk. Be sure to balance the lower financial cost with offsetting costs in terms of risk, availability, or performance.

In many parts of the world, folks are using a blind approach to moving to cloud offerings in the *hope* that it will reduce costs. You can evaluate whether a cloud offering will actually save you money if you identify, measure, and map the costs. Start with a baseline TCO that includes costs like those shown in the first column of the following table. Next, identify the costs you’ll continue to incur and the costs you’ll transfer to the cloud provider (the second and third columns of the table).

Moving Some Storage Functions to Private Cloud Services

<i>Current TCO Baseline Costs (\$A)</i>	<i>Costs You Retain (\$B)</i>	<i>Costs You Move (\$C)</i>
Depreciation		Depreciation
Maintenance		Maintenance
Transformation	Transformation	
Migration		Migration
Microcode updates		Microcode updates

Current TCO Baseline Costs (\$A)	Costs You Retain (\$B)	Costs You Move (\$C)
Administrative staff	Administrative staff	
Power	Power	
Cooling	Cooling	
Floor space	Floor space	
Cost of growth		Cost of growth
Backup	Backup	
SAN infrastructure	SAN infrastructure	
Total Current TCO Baseline Costs (\$A)	New TCO Baseline Costs (\$B)	New Transferred Costs (\$C)

In this example, it's entirely possible that $\$A < \$B + \$C$, and the cloud offering would therefore *increase* total costs. To be sure, cloud providers can, with economies of scale, provide better unit costs. However, with margin and transformation cost and risk, total unit costs might not change as much as you'd expect.

Note: Before jumping on a new cloud offering, be sure to understand your

costs before and after the cloud migration. If total costs don't go down, then you need to be satisfied with the other qualitative benefits that cloud services can offer. And, don't let the cloud providers tell you which costs will come down. You need to have them prove that the costs relevant to you will indeed come down.

Improving Capacity Utilization and Performance with Thin Provisioning

In a traditional storage environment, provisioning is one of the IT department's nightmares — and a fairly high percentage of IT labor costs stem from the time needed to manage and provision storage capacity. If you don't provision enough capacity, users scream at you and business could actually come to a screeching halt.

On the other hand, if you provision too much capacity, you set aside and pay for more capacity than you actually need. This contributes to low capacity utilization rates and inefficient environments. Project this out as data centers grow by 25-30 percent, and you have an expensive proposition on your hands — one that will eventually gain management attention.

You can use *thin provisioning* to solve this problem. Thin provisioning is a software-driven technique that allocates and uses space when it's needed — and *only* when it's needed. Thin provisioning enables you to easily satisfy application requirements without doling out more space than is needed — essentially providing capacity “just in time” as opposed to “just in case.” For example, when you create a 2TB volume using conventional provisioning, you allocate the entire amount even if you have only 300GB of data to store on the volume. Thin provisioning, in that same scenario, only uses the 300GB you actually need, leaving 1.7TB available for other applications as new needs arise.

Thin provisioning lowers TCO because it:

- ✔ Defers capacity purchases by reclaiming allocated but unused capacity and making it available to other applications.
- ✔ Improves application performance if you deploy data dispersion (or wide-striping) technologies.
- ✔ Improves the IT department's responsiveness to new storage demands without downtime.
- ✔ Aligns storage usage with actual business usage.

While many vendors offer thin provisioning, some specific types of thin provisioning go much further than TCO reduction and capacity reclamation. As noted earlier, some types of thin provisioning can also significantly improve storage performance using data dispersion (wide-striping), which allows portions of the data to be stored across many physical drives — as opposed to on one drive — significantly improving application access time.

Aligning Data with Its Business Value

If you're reading this book, your company probably isn't just starting to do business — instead, you've likely been in business for years and you've accumulated a wide variety of storage from several different vendors. And, you've already noticed that some storage is less expensive to use than other storage. In fact, much of the labor in the IT department goes toward trying to set up your storage so that you use it in the most cost-efficient way. Ideally, to lower the total cost of storage ownership, you want to put enterprise applications on high-end, costly disk such as SSD or SCSI (SAS) and applications like payroll and human resources on low-end, less expensive disk such as SATA. This approach helps you to align your applications and data with the business value they represent.

Some customers have found the task of using storage efficiently so overwhelming that they tend to use a monolithic storage structure. That is, they put everything on one tier — typically high-end storage purchased at a general rate. This approach implies that you expect a uniform growth rate for all applications and data — regardless of their individual resource requirements or business value. This monolithic approach to storage has a major drawback: Low-value data and archive data enjoy resources far beyond what they need, making them far more expensive to run than necessary.

Enter *tiered storage*. Using tiered storage, you establish higher storage tiers as higher response/higher cost tiers and lower storage tiers as lower response/lower cost tiers.

With some tiered storage solutions, you can have a storage administrator establish parameters that software uses to move data to a specific tier when the administrator deems appropriate. Or, you can use a policy-driven tool that moves data based on simple data characteristics like age or data type or on a more sophisticated disk characteristic such as frequency of I/O operations.

In the case of I/O, certain types of tiering software will move the application and its data up a tier so that it can run on faster drives and provide better service. The higher-priced drives are justified by the need of the application. Similarly, the

software moves applications like payroll that generate fewer I/O operations down a tier, where they run less expensively.

But tiered storage can be so much more than simply storing data on various storage systems. To truly use multiple tiers of storage effectively, you want to create a single pool of storage with integrated but segmented storage tiers that don't depend on the vendor who supplied the physical storage. In this way, data can move painlessly between tiers as needed, with "as needed" being defined as relative to the data's business value, based on both cost and performance, even as the data's business value changes. This approach ensures that data always resides on the cost appropriate type of storage (see Table 5-1).

Table 5-1 Comparing a Single Storage Pool versus Multiple Storage Tiers

<i>Category</i>	<i>Monolithic Storage Tier</i>	<i>Multiple Tiers of Storage</i>
Pools or tiers	One standard for all data storage	Two to five tiers of storage, matched to the applications: <ul style="list-style-type: none"> • Highest tier for critical applications • Middle tier for everyday applications • Lower tier for archive, tape replacement
Cost	One base rate for all storage	Cost rates matched to disk type: <ul style="list-style-type: none"> • Higher cost disk for a limited number of applications • Moderate cost disk for the majority of applications • Low cost disk for archive, backups, and so on
Price erosion	One base rate for all storage	Base rates matched to storage tiers: <ul style="list-style-type: none"> • Moderate erosion for higher tier • High erosion rate for middle tiers • Very high erosion rate for lowest tiers

<i>Category</i>	<i>Monolithic Storage Tier</i>	<i>Multiple Tiers of Storage</i>
Growth rate	A blended growth plan, usually for the worst-case projects	Rates of growth and capacity planned by tier or category of applications: <ul style="list-style-type: none"> • Higher tiers tend to be more stable and therefore slower growth • Middle tiers tend to have higher growth • Lower tiers often have erratic growth



Marrying tiered storage with virtualization enables you to use your storage most effectively, because you can manage heterogeneous external storage hardware assets from multiple vendors and multiple operating systems using one set of management tools and one universal operator console. Again, this approach allows you to extend the useful life of older assets.

Using one virtual pool of storage, you can

- ✔ Move applications and data outside an array — yes, that’s right. You can virtualize arrays of storage regardless of the vendor from whom you bought them, and virtualized arrays inherit the best features of the parent array.
- ✔ Combine, split apart, re-combine, and provision storage assets across multiple tiers of storage “on the fly,” providing different quality of service characteristics per tier for availability, performance, and recovery.
- ✔ Improve both cost and performance by non-disruptively and automatically moving data to the right tier of storage based on business value.
- ✔ Use a single set of software tools to move, migrate, or replicate data between tiers.
- ✔ Simplify viewing, provisioning, managing, and reporting of all storage resources.
- ✔ Reduce operating expenses (OPEX) related to service level agreement management.

- ✓ Improve your ROA by reclaiming high-performance storage and moving the bulk of older data to less expensive media.

Tiered storage combined with virtualization can improve storage utilization and significantly lower costs (see Figure 5-1). In this example, the bulk of data has been migrated from Tier 1 to Tier 4 (the least expensive tier) to reduce total cost of storage.

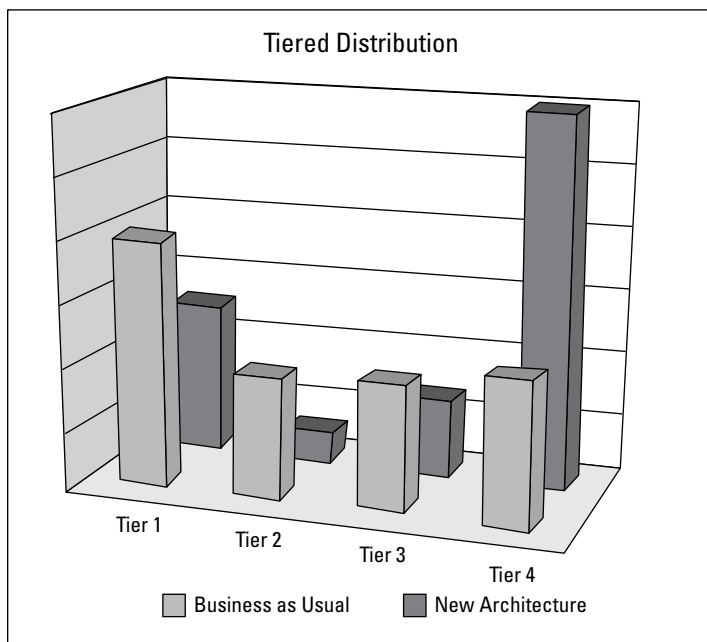


Figure 5-1: You can lower TCO by marrying automated storage tiering with storage virtualization.

Benefits of the “Big Three...”

New technologies to build economically superior storage architectures arrive on the scene all the time. At present, HDS finds that organizations can implement three technologies in

concert to create their individual best storage architecture: storage virtualization, thin provisioning, and automatically tiered storage.

Together, these three key elements can deliver a TCO reduction of 20 percent to 35 percent over older storage architectures. When you use these solutions together as the core of a new storage architecture, the overall impact is greater than the sum of its parts.

Advanced storage architectures can provide a real cost savings when you compare them to traditional, monolithic architectures, as you can see in Figure 5-2.

Table 5-2 identifies the compound effects of using storage virtualization, automatically tiered storage, and thin provisioning.

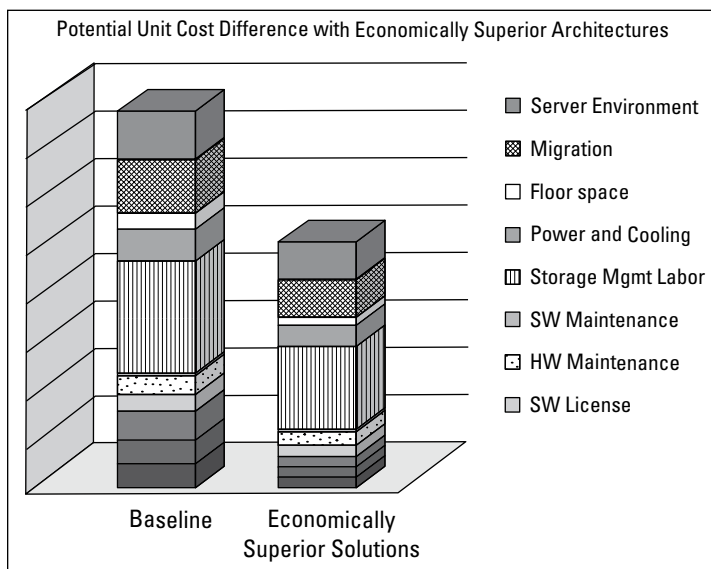


Figure 5-2: Cost reductions of advanced storage architectures.

Table 5-2 Examining the Compounded Effects of the “Big Three” Economically Superior Solutions

<i>Factor</i>	<i>Effect</i>
Significant storage reclamation	Experience both a one-time reclamation and a significantly lower cost of growth over time.
Return on assets (ROA)	Utilize assets better, even beyond their depreciated life.
Total capacity	Do more with less and maintain the same quality of service.
Data migration	Migrate heterogeneous storage without disruptions.
Backup and recovery	Share backup services, replication, snap copy management, and more. Reduce backup and recovery times and reduce the cost of copies with multi-tiered storage.
License costs per terabyte	Reduce costs even as storage capacity grows.
Scheduled downtime	Perform non-disruptive maintenance, migration, upgrades, and provisioning.
Change management	Manage heterogeneous assets, make non-disruptive configuration changes and spend less labor time on provisioning through a common interface.
Environmental	Reduce floor space, power, and cooling as measured per terabyte (kVA/TB or kW/TB).



HDS consultants have formally and informally measured TCO and ROA with clients who have implemented these key technologies and have found that the resulting TCO improvement averages approximately 25 percent within the first year after implementation.

The economic impact of virtualization at Overstock.com

Overstock.com (shortcut O.co) is an online retailer offering a wide variety of high-quality, brand-name merchandise and services at discount prices, including bedding, home decor, appliances, watches, jewelry, electronics, sporting goods, clothing, shoes, cars, travel and insurance. The company gives its customers an opportunity to shop for bargains conveniently, while offering manufacturers, distributors and other retailers an alternative sales channel. The volume of business data and customer records that Overstock.com generates and retains is growing at an aggressive rate that exceeds 40 percent annually, and Overstock.com must be able to store and retrieve data for longer periods of time.

To continue providing reliable and high-performance storage support for expanding production operations and application demands, Overstock.com determined that the storage environment should be nearly limitless in its ability to expand with growth requirements.

Overstock.com developed several goals to help the company improve service levels and data access while reducing operational inefficiencies and related costs:

- ✔ Simplify the storage environment
- ✔ Align storage tiers to business needs

- ✔ Simplify data migration
- ✔ Reduce operational expenditures
- ✔ Reduce capital expenditures

Using the cost categories presented in Chapter 2, Overstock.com identified 14 distinct storage cost categories to reduce. At the beginning of the project, Overstock.com created a TCO baseline for the storage estate as it existed with the old architecture and has used this baseline to compare architecture improvements since then.

Overstock.com chose to implement virtualization, thin provisioning, and automated storage tiering. They chose the HDS platform because the HDS platform offers controller-based virtualization that enables IT staff to manage the entire data center from a single interface as one virtual storage system for both internal storage and externally attached heterogeneous storage, including mainframe storage from multiple vendors.

Overstock.com implemented thin provisioning to provide virtual storage capacity from a storage pool. Because thin provisioning has no operational overhead, it decreases the administration costs normally incurred with provisioning new storage. It also improves the availability of applications by reducing the downtime needed for the storage provisioning process.

(continued)

(continued)

Overstock.com also implemented a virtualized three-tier storage solution that automatically classifies data according to the service level needs of business applications; as needed, data moves seamlessly to the correct storage tier, which improves capacity utilization.

While you can convey the impact of a virtual, thinned, and automatically tiered storage architecture in several impressive ways, a lower unit cost is perhaps the management metric that most matters. IT departments cannot necessarily control the demand for storage, but through strategic and tactical investments, they can drive down unit costs. In the first year after Overstock.com switched architectures, the average unit cost of a GB/year of storage went down by 43 percent. The improvements continued in the second year with a 37 percent reduction in unit costs.

In addition to the quantitative benefits, Overstock.com has experienced several qualitative benefits:

- ✔ Migrating data, a task that took upwards of five hours before changing the storage environment, now takes less than 30 minutes with fewer people.
- ✔ Overstock.com now uses 80 percent or more of its storage in the virtualized environment using thin provisioning.
- ✔ The time Overstock.com needs to plan for and implement new applications has been significantly reduced and, in some cases, eliminated altogether,

improving time to market for the new applications.

- ✔ Based upon business needs, Overstock.com has been able to monitor performance of key applications and seamlessly migrate data to higher or lower tiers.
- ✔ Operational efficiencies have improved. Storage administrators have regained 40 to 50 hours each quarter because they no longer need to monitor the performance of every application, watching for hot spots that required moving data to another location. In addition, provisioning tasks that typically took two to three hours the traditional way now take less than 30 minutes.

“We depend heavily on the Hitachi Data Systems platform to reliably run our business and scale as we grow,” says Carter Lee, VP, Technology Operations at Overstock.com. “The Hitachi platform allows us to seamlessly virtualize and manage all the storage as one logical array. This architecture simplifies management, allows us to scale more cost-effectively, and gives us the flexibility that we need to keep up with the demands of the business. By virtualizing external storage assets, Overstock.com was able to simplify the entire technology refresh process. Migrations that typically took upwards of five hours took less than 30 minutes with the virtualization solutions.”

Chapter 6

Show Me the Money

In This Chapter

- ▶ Exploring the flip side of reducing operating costs: generating profits
 - ▶ Examining the characteristics of the Storage Economics hero
-

Reducing operating costs is a hot topic, but the flip side of reducing costs is increasing and generating new types of “profit” from the right storage architecture. Many IT shops are implementing “chargeback” strategies to appropriately bill out storage resources to those internal businesses that use them. Granted, most IT departments are not profit centers, and some do not desire to be, but you can apply some basic concepts no matter your designation as a cost center or profit center. Sometimes, to be efficient, you can think like a profit center to change the way you allocate and manage storage resources. This chapter explores some of these concepts and also outlines the characteristics of Storage Economics heroes — those who learn the language that financial folks speak and apply it to storage investments.

Reducing Costs Isn't the Whole Picture

Many businesses that are thought to be efficient have implemented “chargeback” strategies and are increasingly transforming their IT cost centers into IT profit centers. With a chargeback scheme in place, internal IT departments literally bill adjacent business units for the storage they consume, essentially holding them accountable for their share of storage consumption and discouraging waste and capacity overconsumption.

If you are considering or beginning to implement a chargeback strategy, you need to consider some foundational technologies and best practices. Using a virtual and consolidated storage architecture will immediately reduce your management efforts and decrease your software costs by changing the calculation rates for storage management tasks, such as load balancing, data recovery protection, and more. These technologies will also allow you to offer a more flexible storage service catalog that incorporates fast provisioning, multiple tiers, archive tiers, and different data protection options to enhance your storage portfolio offering and attract new clients.

Your IT bosses might not ask you to “show them the money,” but you can make architectural alterations to increase storage revenue for billable services. For example, you can:

- ✔ Offer different levels of \$/TB, each with varied margins, with automatically tiered storage. This approach encourages consumers to use lower tiers and discourages them from using higher tiers. You also can charge for data movement between tiers — perhaps a one-time charge of \$/TB moved, like a brokerage services charge for buying or selling financial instruments.
- ✔ Offer a lower-cost archive tier with different price and performance rates and different margins of profit. You also can offer discovery and index-search capabilities for faster discovery (important for legal and medical applications).
- ✔ Offer Tier 0 capability that uses solid state drives, in addition to a lower cost archive tier. Usage might be small at first, but if tightly integrated with tiered allocation, Tier 0 can have a good profit margin and extend the levels of storage service you offer.
- ✔ Offer different target tiers of storage — at different rates and margins — for data recovery protection.
- ✔ Use capacity-on-demand or policy-driven provisioning, letting applications or hosts request new capacity and keeping humans out of the loop. You measure the consumption levels and charge the consumer only for allocated or consumed storage.
- ✔ Offer and charge for de-duplication services to help your customer use less space.
- ✔ Offer different backup options to meet user requirements and charge different prices for different recovery time

objective and recovery point objective mediums — that is, disk versus tape backup options.

The Storage Economics Hero

It isn't enough to know technology anymore; you need to become a storage economist — someone who can apply technology benefits to cost reduction and understand the implications. Navigating difficult economic and budget sessions in the meeting rooms of your company often requires a new set of skills.

First, think like an economist

Look at the macroeconomic (big picture) situation and your microeconomic situation — the economics of your company — to determine whether your company might be heading into an austerity phase. Understand your capital approval process and slip in an un-forecasted request to invest in technology and practices that can get your storage capacity demand through another drought. In your justification, use the following phrases to capture the attention of capital owners and make a lasting impression:

- ✔ I want to improve ROA and invest in some areas to increase the usage and value of what we already have. In fact, with a small investment I can reclaim enough capacity to meet our growth needs by ___ months (fill in the blank).
- ✔ I want to sweat our assets without incurring additional new costs, but that will take some small investments now.
- ✔ By changing or influencing consumption behaviors, I can squeeze more efficiency from our IT infrastructure and work to continue storage optimization.

Did you catch the key words in the preceding list? *ROA, reclaim, sweat, optimize*. Keep in mind that these are just a few of the many key words economic heroes like to use.

Second, talk like an accountant

You need to understand your costs, all the costs — not just purchase price. Chapter 2 outlines 34 types of storage costs and Chapter 3 helps you to isolate and prioritize your costs

associated with storage. Focus on your high priority costs and implement tactical plans to invest and reduce these costs now. Measure your costs before and after you act; that way your tactical technical decisions will make you an economic rock star.

Finally, act like a technologist

In times of a capital crunch, all companies reduce budgets, and these reductions are more painful in the storage area where demand and growth continue without respect to budget constraints. Take action to invest and extend what you have:

- ✔ **Consolidate.** Identify and act to collapse multiple, smaller storage islands into a single pool of manageable storage. Virtualization can be a key enabler to make this happen with limited time, effort, and business impact.
- ✔ **Virtualize.** Storage consolidation is a key tool for reducing capital costs, but it also can increase operational costs since it increases the number of applications and users affected when changes are necessary for the storage system. When you consolidate 100 applications onto a common storage system, you cannot stop all 100 applications while you make a change to the storage system. Instead, you need to schedule this change with all the application users. Storage virtualization masks these changes and reduces operational costs. Virtualization also helps to map new cost-saving technologies to existing storage assets and provide a measurable ROA.
- ✔ **Reclaim.** Don't delete data; instead, reclaim unused space using tools like zero page reclaim, virtualization, and active archive.
- ✔ **Optimize.** Use tools like thin provisioning, de-duplication, and policy-based management and provisioning to optimize and re-balance existing capacity.
- ✔ **Oversee.** Start capturing and reporting on usage abuse — those applications or departments that notoriously under-utilize or lock-up precious capacities.

In sum, the Storage Economics hero calculates the total cost of ownership (TCO), identifies important costs to reduce, measures them, implements changes — ones that align cost-reduction services to identified costs — and measures again to verify progress.

Chapter 7

Ten Ideas for Reducing Total Cost of Storage Ownership

In This Chapter

- ▶ Understanding the best practices for implementing a storage cost reduction strategy
-

Your company's economics are tied to global and national economics, and the economics of your company impact the economics in the IT department. This is the nature of trickle-down economics. Storage Economics uses proven and stable principles that can help you navigate economic trouble spots that call for cost cuts.

The previous chapters walk you through how to identify, map, define, and measure storage costs. This chapter helps you move forward by tying these chapters together and outlining some important best practices for technology implementation and cost reduction

Consider these ten ideas if you're looking for actionable ways to kick off a new cost-reduction strategy in your IT department:

- ✔ Price is not the same as cost, so focusing on price reductions alone will have very little impact on reducing your costs. For example, if the total cost of owning disk is around 15-17 percent of TCO, negotiating a 10 percent reduction in disk price will only impact TCO by 1.5-1.7 percent. Focus on the right technologies to reduce the costs important to your environment.

- ✔ Focus on cost-reduction efforts in a down economy by pulling a fine-tooth comb over operational costs. Review labor, power, cooling, maintenance, data protection, and data recovery protection to determine where you can trim costs.
- ✔ Save some ammunition for later. Plan to implement cost-reduction items over two years at a minimum, and maybe longer.
- ✔ Do the easy things that impact total cost first, but don't do all of them. Save some easy cost-reduction options for later, too. Also save the harder ones — whether political or organizational — for later.
- ✔ Isolate and work to reduce costs with those people who measure or care about the costs and are responsible for them. You need allies in cost-reduction planning, and the IT department may not be directly responsible for all the related costs.
- ✔ Beware of quick-fix solutions that appear too good to be true. Make sure that vendors are helping you reduce the costs that are important to you — not the costs their solutions can reduce for you. And don't forget: Most solutions to reduce costs are not free.
- ✔ Attitudinal and political changes are the most difficult to implement.
- ✔ Changes often require rewards and penalties, so factor that into your plans.
- ✔ Understand the differences between direct (hard) costs and indirect (soft) costs.
- ✔ Measure twice, cut once. That is, make sure that you can measure the costs you can impact now and then measure those same costs again later.

Use Storage Economics principles to identify, measure, and reduce the total cost of storage!

Storage Economics helps determine the real cost of the storage infrastructure by recognizing that the cost of acquiring assets is just a small percentage of their total cost to own. Storage Economics focuses on measuring the total cost of ownership (TCO) to align information technology investments with corporate goals.

- **Understand why price doesn't equal cost** — see how the acquisition cost is only a small percentage of the total cost of ownership
- **Learn what costs comprise total cost of storage ownership** — get informed about the costs related to storage
- **See how proven solutions improve performance and reduce costs** — know the benefits of economically superior solutions
- **Discover what deployment strategies can be used** — examine a case study to see how storage economics works in the real world
- **Find out how to evaluate storage system proposals** — learn what key information to provide to assure that vendor proposals focus on costs important to you

Hitachi Data Systems provides best-in-class information technologies, services, and solutions that deliver compelling ROI, return on assets, and demonstrable results. Hitachi Data Systems improves IT costs and agility with solutions that are virtualized, automated, cloud-ready, and sustainable. Hitachi Data Systems believes that data drives our world — and information is the new currency. To learn more, visit www.HDS.com.



Open the book and find:

- Why traditional approaches do not accurately measure the cost of owning storage
- What economically superior solutions can do to help you reduce the total cost of owning storage
- Ways to match costs you want to reduce to solutions that reduce them
- How you can become an economic hero in your organization

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

For Dummies®
A Branded Imprint of



ISBN: 978-1-118-17787-7
Book not for resale