

**Making Everything Easier!™**

**Oracle Special Edition**

# **Storage Tiering**

FOR  
**DUMMIES®**

**Learn to:**

- Implement tiered storage in your organization
- Intelligently improve storage performance and increase capacity
- Drive down storage costs

*Brought to you by*

**ORACLE®**

**Michael Wessler, OCP & CISSP**





***Storage  
Tiering***  
FOR  
**DUMMIES®**  
ORACLE SPECIAL EDITION

**by Michael Wessler, OCP & CISSP**



Wiley Publishing, Inc.

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

## Storage Tiering For Dummies®, Oracle Special Edition

Published by

**Wiley Publishing, Inc.**

111 River Street

Hoboken, NJ 07030-5774

[www.wiley.com](http://www.wiley.com)

Copyright © 2011 by Wiley Publishing, Inc., Indianapolis, Indiana

Published by Wiley Publishing, Inc., Indianapolis, Indiana

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

**Trademarks:** Wiley, the Wiley Publishing logo, For Dummies, the Dummies Man logo, A Reference for the Rest of Us!, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. Oracle is a registered trademark of Oracle International Corporation. All other trademarks are the property of their respective owners. Wiley Publishing, Inc., is not associated with any product or vendor mentioned in this book.

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

ISBN: 978-1-118-06262-3

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1



WILEY

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

# Contents at a Glance

---

<b><i>Introduction</i></b> .....	<b>1</b>
About This Book.....	1
Icons Used in This Book.....	2
<b>Chapter 1: Today's Storage Challenges</b> .....	<b>3</b>
The Increasing Demand for Storage .....	4
Data Access Rates: I/O Keeping Up with CPU .....	6
Increasing Data Longevity.....	6
Proliferation of Storage Alternatives .....	8
Keeping Storage Costs Manageable.....	11
The Need for a Better Solution.....	12
<b>Chapter 2: What Is Storage Tiering?</b> .....	<b>13</b>
Defining Storage Tiering.....	13
Data Usage Classification.....	15
Tiers of Storage .....	16
Tiering Methodologies .....	18
Balancing Storage Cost, Speed, and Capacity .....	20
<b>Chapter 3: Oracle's Complete Tiered Storage Solutions</b> .....	<b>25</b>
Complete Coverage of Tiered Storage Needs.....	26
Correct Storage for Each Level of Tiering.....	26
Integrating Multiple Storage Tiers .....	30
<b>Chapter 4: Implementing Oracle Tiered Storage</b> .....	<b>35</b>
Environments That Benefit	
the Most from Tiered Storage .....	36
Oracle Tiered Storage Use Cases .....	37
Selecting the Product Stack.....	40
Installation and Configuration.....	41
Support and Operations.....	42
<b>Chapter 5: Ten (Okay, Five) Things to Look For in a Tiered Storage Solution</b> .....	<b>43</b>
Are You a Good Fit? .....	43
Do You Know Where to Start? .....	44
Are You Using All the Options?.....	44
Are You Using Software Wisely? .....	44
Are You Avoiding Common Barriers? .....	44

## **Publisher's Acknowledgments**

We're proud of this book and of the people who worked on it. 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](mailto:info@dummies.biz). For details on licensing the *For Dummies* brand for products or services, contact [BrandedRights&Licenses@Wiley.com](mailto:BrandedRights&Licenses@Wiley.com).

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

### ***Acquisitions, Editorial, and Media Development***

#### **Senior Project Editor:**

Zoë Wykes

**Editorial Manager:** Rev Mengle

#### **Senior Business Development**

#### **Representative:**

Karen L. Hattan

#### **Custom Publishing Project**

**Specialist:** Michael Sullivan

### ***Composition Services***

#### **Senior Project Coordinator:**

Kristie Rees

#### **Layout and Graphics:**

Carrie A. Cesavice,

Christin Swinford

#### ***Special Help from Oracle:***

Michael Brown, Amy Thompson,

Aaron Newcomb, Kerstin Woods,

Suzanne Blackstock

---

## **Publishing and Editorial for Technology Dummies**

**Richard Swadley**, Vice President and Executive Group Publisher

**Andy Cummings**, Vice President and Publisher

**Mary Bednarek**, Executive Director, Acquisitions

**Mary C. Corder**, Editorial Director

## **Publishing and Editorial for Consumer Dummies**

**Diane Graves Steele**, Vice President and Publisher, Consumer Dummies

## **Composition Services**

**Debbie Stailey**, Director of Composition Services

## **Business Development**

**Lisa Coleman**, Director, New Market and Brand Development

# Introduction



**D**ata growth in enterprises and organizations has exploded due to new data sources, the ways we use data, and the business and regulatory requirements to retain data for years. Large unstructured data objects such as images, documents, and video are now as necessary as raw numbers in database tables.

Current IT storage techniques have an increasingly difficult time supporting the new capacity requirements and keeping storage costs under control; a new approach is needed. That's where tiered storage solutions enter in, because tiering your storage can meet these massive growth requirements without breaking your IT budget.

## *About This Book*

This book consists of five short chapters, each written as a stand-alone chapter, so feel free to start reading anywhere and skip around throughout the book!

**Chapter 1: Today's Storage Challenges.** We look at why storage requirements are growing so fast. Then we look at the technologies and techniques currently in use and why they aren't sufficient for the future.

**Chapter 2: What Is Storage Tiering?** Here, we get into the nuts and bolts of different types of data storage, how they work, and why tiered storage is a better solution.

### **Chapter 3: Oracle's Complete Tiered Storage Solutions.**

This chapter looks at Oracle's wide selection of hardware and software tiered storage solutions.

**Chapter 4: Implementing Oracle Tiered Storage.** We delve into how Oracle tiered storage is implemented in the real world and what you need to consider for your implementation.

**Chapter 5: Ten (Okay, Five) Things to Look For in a Tiered Storage Solution.** Here, in that famous *For Dummies* style, we give you the "Part of Tens" detailing items to help get you started with implementing a tiered storage architecture.

## *Icons Used in This Book*

Throughout this book, we occasionally use icons to call attention to important information that is particularly worth noting. Here's what to expect.



If you see an icon that says remember, you may want to, uh, remember it. We won't have deep meaning-of-life stuff, but it may be good to know for later.



Sometimes you just have to know a term or technical details to understand a larger topic. Or, it could just be that we want to throw in cool terms to try to impress people. Either way, this icon identifies techie stuff.



This icon usually denotes something the author wishes someone had told him before he learned it the hard way! Keep these items in mind to make life easier.

# Chapter 1

---

## Today's Storage Challenges

.....

### *In This Chapter*

- ▶ Taking a look at the rising demand for storage
  - ▶ Evaluating data access rates
  - ▶ Ramping up data longevity
  - ▶ Understanding storage alternatives
  - ▶ Holding storage costs in check
  - ▶ Finding a storage solution
- .....

**I**ntelligent management of an organization's storage is an increasingly difficult, yet required task for today's IT staff and leaders. Gone are the days of simply buying more disk when you are about to run out. Careful analysis of both the technical and business aspects as part of an enterprise storage plan is how smart businesses are implementing storage today. If you do things right, you will save money, improve system performance, and meet diverse and growing storage requirements.

In this chapter, we look at the storage challenges faced by organizations so that you can make the right decisions for your organization.

## *The Increasing Demand for Storage*

No doubt about it, the amount of data being stored today is increasing at an alarming rate. People everywhere keep digital forms of high definition images, detailed maps, documents, music, video, and other large objects in addition to the raw numbers and text they're used to storing.

Need a quick example? Take a look at your smartphone and all the music and images on it that you have backed up (hopefully) to your PC; and that's just the tip of the iceberg. If you need instructions, just check YouTube — which is another example in itself!

Far beyond just personal use, storage increases are impacting business, scientific, and government agencies as well. Here are some examples:

- ✔ The number of online sales transactions that happen every day on the Internet. Online retail sales in the U.S. alone are expected to grow to \$248.7 billion by 2014 according to one Forrester Research sales forecast. Also consider the treasure trove of demographic data associated with those sales.
- ✔ Sales at global retail companies such as Wal-Mart, which alone generate more than 1-million customer transactions per hour. Add to that the supply chain management ramifications.

- ✓ Detailed medical images captured by high resolution imaging machines. Regulatory requirements may dictate retaining this data for extended periods of time, and best practices are that this period is longer than the patient's life.
- ✓ High resolution global maps and GPS information used by government, military, and private individuals.
- ✓ Content rich Web 2.0. Although personal use of such applications as Facebook and YouTube are obvious, smart businesses are leveraging similar techniques from their own portals to build user communities through the use of social networking and on-demand video.

Industry experts are unified in the belief that data storage demands are increasing; the only question is by how much. According to the February 27, 2010 issue of *The Economist*, data is growing at a compound annual 60 percent rate. A 2008 study by IDC (International Data Corporation) stated that in 2010 alone 1200 exabytes would be generated. **Note:** For perspective on how big an exabyte is, consider that one terabyte is 1024 gigabytes and 1024 terabytes is one petabyte and 1024 petabyte is one exabyte!

Exact numbers are impossible to determine and growth estimates will vary between sources. However, the obvious direction is enormous growth in storage requirements for both structured data (databases) and unstructured data (images, video, sound, documents).

## ***Data Access Rates: I/O Keeping Up with CPU***

Storage access rates have failed to keep up with the remarkable increases in computer processing speeds. While central processing unit (CPU) speeds have skyrocketed, storage access speeds have merely increased. Although the computer can process data quickly, there *is* a disconnect because it must wait on disk Input/Output (I/O) latency. Despite the increased speed of the CPUs, the overall system is still no faster than the storage it runs on.

Disk storage is already at a disadvantage because it is mechanical, making it inherently slower than other components in the system, such as memory and CPU speeds. The problem is exacerbated if the storage solution suffers from poor design or implementation.

In terms of cost, consider a few seconds to several minutes of delay per computer action, multiplied by multiple iterations per day, again multiplied by the number of users, all of which equates to tangible wasted time and therefore wasted money for your organization. At an executive level, the business leader making key decisions isn't getting the necessary information when it's needed — which can be disastrous.

## ***Increasing Data Longevity***

Not only do you now have to keep more data, but you also need to keep it longer. It used to be that after data was no longer operationally needed, you could delete

it. Today, two key factors have caused an increase in data longevity: regulatory requirements and trend analysis.

## ***Regulatory requirements***

Legal and government regulations have specific requirements for what data you keep, in which format, and for how long, including various retention periods for key data elements. These requirements vary by industry.

Key data retention examples such as Sarbanes-Oxley's requirement to keep audit data for seven years is well known, but other industries often require between three to six years for data elements. At the extreme end, hazardous materials may require 30-year retention periods and other periods may be indefinite.



Some acts such as the Health Insurance Portability and Accounting Act (HIPAA) and Sarbanes-Oxley have always been known for stringent data requirements. A more recent trend is the push towards encryption of Personally Identifiable Information (PII). Security concerns over compromised data are driving the need to encrypt key data elements. If your organization hasn't reviewed what data it encrypts, it should.

## ***Trend analysis***

By knowing how your industry performs over time, you can make smart business decisions to gain a competitive advantage. To do this, you need to retain the detailed operational data to perform trend analysis.

Currently, operational data is often rolled off to a data warehouse or datamart to perform trend analysis with Business Intelligence (BI) tools. Business leaders use these BI tools to identify market trends and business opportunities.

Non-business uses of BI tools include clinical studies, environmental studies, and other scientific research. However, to use BI tools, the underlying data must be present and accessible. Because deleting/purging data would inhibit BI tools' capabilities, you must retain the data — but you can move it to long-term storage as long as it remains accessible.

## ***Proliferation of Storage Alternatives***

Organizations often have a wide mix of storage mediums in use. Historically, internal disks were used for active systems and tape was used for backups. While this configuration was a necessary “first step,” technical advancements and business needs rendered this configuration unsuitable.

As time and technology moved along, more options became available in the form of attached storage devices such as Storage Area Network (SAN) and Network Attached Storage (NAS). These devices allowed capacity to rapidly grow, enabled data sharing, and included more advanced management software to maintain it.



Storage Area Network (SAN) and Network Attached Storage (NAS) are large storage devices containing multiple disks. These disks are formatted into usable allocations and presented to attached servers to be used instead of/or to supplement internal storage. SANs and NAS can be large and complex, requiring specialized software and skill sets to maintain them.

A more modern approach is to add flash and unified storage into the existing environment of attached storage. Flash and unified storage are complementary technologies to traditional disks that offer different performance, price, and functionality profiles.

For example, flash storage costs more than disk on a capacity basis, but it sustains so many more input/output operations per second (IOPS) that it is often less expensive in IOPS intensive environments. Similarly, unified storage solutions combine NAS and SAN capabilities into a single platform, which offers somewhat lower IOPS than traditional Fibre Channel SAN devices but with much lower cost on a capacity basis and simplified deployment and management — a growing cost in today's datacenters.



Flash is high-performance, non-volatile memory that is able to store and recall data with much lower latency than disks. Flash storage can be leveraged to expand server cache or utilized at the storage layer as a disk replacement or to reduce the latency of a disk-only system.

See Table 1-1 for a summary of the storage media and their strengths and weaknesses.

**Table 1-1 Storage Media and Attributes**

	<i>Storage Type</i>			
	<i>Flash</i>	<i>Fibre Channel Arrays</i>	<i>Unified Storage</i>	<i>Tape</i>
<b>Capacity</b>	Low	High	High	Very High
<b>Capacity/\$</b>	Low	Medium	High	Very High
<b>IOPS</b>	Very High	High	Medium to High	Low
<b>IOPS/\$</b>	High	Medium	High	Low
<b>Bandwidth</b>	Very High	High	High	Medium
<b>Bandwidth/\$</b>	High	High	Very High	Medium

As you see in the table, each media type has different strengths and weaknesses. No one solution meets all the needs for an organization.



The technical study of storage types is ever-changing and can be a profession in itself. MLC/SLC Flash, FC SAN, iSCSI, FCoE, InfiniBand, and NAS subtypes are just some of the technologies associated with storage. Software that simplifies the management of these complex technologies is well worth the cost.

## *Keeping Storage Costs Manageable*

Data storage requirements are growing, but IT budgets are often unable to keep pace. Be it the notorious “unfunded government mandate” or simply a new business need to store large data types, the purchase of larger and faster storage devices is occurring.

Beyond the purchase price of the storage medium, you also need to consider these additional costs:

- ✓ Floor space for large storage devices
- ✓ Electrical costs for additional storage devices
- ✓ Cooling costs for the data center
- ✓ Network bandwidth to move data in large, distributed environments
- ✓ Management software for storage
- ✓ Administration costs for people to manage the storage
- ✓ Corresponding development and test environments
- ✓ Backup and recovery of the data
- ✓ Disaster Recovery (DR) involving the data

Smart IT staffs can mitigate some of the increasing costs. Examples include smaller-sized development and test environments. Standardizing and automating storage procedures are another common practice in efficient IT shops. Administration costs can be expensive so simplifying environments and leveraging management tools can reduce administrator costs.

## *The Need for a Better Solution*

Simply buying more of the same disk to meet your rapidly increasing data needs is feasible as long as you have unlimited funding, unlimited energy, and unlimited datacenter floor space; unfortunately, the vast majority of us are not in that situation.

A better and smarter solution is to review how the data you have is used and retained. Not all data has the same usage characteristics and not all data warrants being on the fastest disk or flash available. Knowing your data characteristics gives you options in managing that data.

Organizing your storage needs into tiers based on access speed, capacity, and retention requirements allows you to use lower cost storage media options appropriate to the data.

It is this *storage tiering* that will allow you to reduce storage costs and management overhead while increasing capacity and performance, and meeting more stringent service level agreements.

## Chapter 2

# What Is Storage Tiering?

---

### *In This Chapter*

- ▶ Explaining storage tiering's theory and goals
  - ▶ Categorizing data
  - ▶ Identifying storage tiers and methodologies
  - ▶ Taking in storage cost, speed, and capacity
- 

**T**iered storage solutions use different storage technologies and configurations to store data based on the data's criticality and usage frequency. By using more expensive, fast storage tiers for your most critical data and less expensive storage tiers for the large amount of less used data, your organization can realize substantial storage savings.

In this chapter, we look at the details of how storage technologies make up different storage tiers, how data is classified and assigned to a tier, and the cost benefits of tiered storage.

## *Defining Storage Tiering*

*Storage tiering* is the technique of placing data on different storage technology or different configurations of the same technology (tiers) based on the characteristics

of that data. Essentially, several different storage technologies are used to store data instead of providing the same performance, reliability, and cost for all data. By intelligently allocating storage tiers, high performance access to critical data is achieved while lowering the total cost of storage ownership.

In tiered storage, different storage technologies and configurations (flash, disk, tape) are allocated for high-use, less-accessed, and seldom-accessed data; these form *tiers*. Generally, the fastest storage technology will be the tier for the most frequently used data.

Remaining tiers will be defined based on their performance access requirements and the value associated with accessing them, with the slowest tier reserved for the least used data.

Each data element to be stored is categorized based on how often it is accessed and the value associated with that access. The data is then stored on the corresponding storage tier defined for its data classification (that is, critical, frequently accessed data on the fastest storage). This classification is dynamic based on data access patterns, age, and other factors.



Lowering total storage cost is the key driver for storage tiering. In theory you could store all your data on the fastest storage available, but that would be cost prohibitive as data grows. Knowing that faster disk or flash drives are more expensive than slower disk or tape, you purchase a small amount of fast storage for your most critical data and put the less critical data on cheaper storage.

## Data Usage Classification

Not all data is the same, especially in terms of its business value or how often it is accessed. These differences in data are critical to understand because they define how your data is tiered.

A small amount of an organization's data will be in high demand and critical because it is accessed frequently. The high demand data is critical because it is either new and the subject of current work (such as recent sales figures) or because it has high usage characteristics (such as a database index or lookup table).



Eighty percent of data is never used again after 90 days according to a 2010 end-user survey by the Association for Information and Image Management (AIIM). Therefore, why treat all your data the same when it's really not?

The remaining data varies between regularly used, seldom used, and finally data that is simply kept to meet regulatory requirements. This data accounts for the vast majority of data stored in an organization.

Table 2-1 shows further refinement of data classification.

<b>Category</b>	<b>% of Data</b>	<b>Examples</b>
I/O Intensive	3–5	Indexes, lookup tables
Mission Critical	10–20	Web-based sales application
Support	20–25	Test systems, reference data
Archive	45–60	E-mail archives, legal records

As the table shows, only a small percentage of data is used with heavy frequency. A larger percentage is used less frequently, but is necessary to support operations (mission critical and support). The largest percentage (and this percentage is growing) is the archive data that is very seldom (if ever) used, but is necessary to retain.

Data's classification often changes over time as it progresses through its lifecycle. Current sales figures may fall into the I/O intensive category but as the data ages, it may become support or archive data.

## ***Tiers of Storage***

Physical storage technologies such as flash, disk, and tape have different characteristics. Access speed, capacity, longevity, moving parts, electricity consumption, and cost are common characteristics to describe and compare storage mediums.

Within storage tiering, it is necessary to define the characteristics of each tier. Classification criteria are speed and capacity with an average cost per gigabyte assigned. Once the available storage is classified into tiers, then data can be assigned to those tiers.

### ***High-performance flash storage***

The fastest and most expensive storage medium (on a cost per TByte basis) is *flash storage*. Flash is nonvolatile memory used to store smaller amounts of data, but is extremely fast to access, and requires less power. Thumb drives and Solid State Drives (SSDs) are both examples of the underlying flash technology. A newer technology, flash storage is extremely fast but is too

expensive to be adopted as the single storage tier for an organization.

Flash storage is suitable for your most critical, highly used I/O intensive data accounting for approximately 3–5 percent of your total storage. It is also an inexpensive method of achieving high performance when considering cost per I/O Operations Per Second (IOPS).

### ***High-performance disk***

High-performance disk storage is fast and offers far greater capacity than flash. This is the fastest disk available, is configured in a RAID configuration to promote high performance, and may be attached externally to the server or through a shared SAN using high-performance Fibre Channel (FC) attached storage arrays.

As data moves off flash, it still needs to be accessed frequently, but not as frequently as previously. A primary storage-performance disk is the correct tier for mission-critical data that requires fast access, but capacity requirements have exceeded the capabilities of flash storage. On average, it accounts for 10–20 percent of your total storage and can be made up of either Fibre Channel arrays or high-performance unified storage configurations.

### ***High-capacity disk***

A disk with slower access speeds, perhaps using a RAID configuration that promotes capacity over speed, is used for a secondary level of disk storage. This storage is not as fast as a performance disk, but has greater capacity and is less expensive per TByte of capacity. Serial Attached SCSI (SAS) and Serial Advanced

Technology Attachment (SATA) drives are examples of lower performance, high-capacity disk.

Data suited to a secondary storage capacity disk tier is support or infrastructure data. It is data that is accessed infrequently, is large in size composing 20–25 percent of your total storage, and is typically comprised of high-capacity unified storage configurations.

### *Long-term storage*

Tape storage is ideal for long-term storage. Although it has the slowest access speed of the storage tiers, tape also has the highest capacity and lowest cost of ownership.

Long-term storage accounts for approximately 45–60 percent of total storage, and this percentage is increasing. Table 2-2, gives you a comparison of storage tiers with corresponding data classifications.

<i>Data Classification</i>	<i>Medium</i>	<i>Capacity</i>	<i>Cost</i>
I/O Intensive	Flash	3–5%	\$40–\$54/GB
Mission Critical	Performance Disk	10–20%	\$7–\$22/GB
Support	Capacity Disk	20–25%	\$1–\$6/GB
Archive	Tape	45–60%	\$0.25–\$1/GB

As you can see, the smaller percentage of highly critical data is considerably more expensive than the slower, greater-capacity, and lower-cost media.

## ***Tiering Methodologies***

How data is classified into tiers can vary based on your implementation. Different technologies can determine data's classification and corresponding storage tier. Furthermore, as data changes classification, the change needs to be recognized and the data migrated to the appropriate storage tier.

There are different methodologies to tier data, and each one corresponds to a general technology. The differentiating factor is whether the application is aware of the tiering requirement and makes tiering decisions, or it is unaware of storage tiering.

### ***Application-aware tiering***

The best example of application-aware tiering is a database configured to partition data. In this configuration, the database is the application that is tiering-aware. Data within a database is stored in tables that are logical structures. Each table is assigned a physical structure, which is one or more data files located on a storage medium. As data is created, it is inserted into the logical table that maps to data files on a storage device.

Partitioning allows tables to be divided up into smaller, portable table segments based on specific data characteristics. A common characteristic on which to partition data is the creation date, which accurately reflects the data's age.

Storage tiering is achieved by placing new data partitions on flash or high-performance disk tiers. Older data partitions are placed on secondary storage

capacity disks or long-term storage devices. Movement of the partitions from faster to secondary storage capacity would be scripted by the Database Administrator.

Management of the storage tiers is performed by the database software; thus it is application aware.



See Chapter 4 for detailed information on application-aware tiering with Oracle databases.

### ***Application-transparent tiering***

Storage tiering not explicitly implemented by the application is said to be application transparent. In these cases, storage tiering still occurs, but the application is unaware of tiering and does not manage the data classification or tiering processes.

Management of data files, blocks, bytes, and the reading and writing of those objects to a physical storage medium such as disk is managed by the *file system*. When an application issues an I/O (Input/Output) request to read or write data to disk, it is the file system that implements the request on behalf of the application.

File systems with specialized software automatically categorize data into a tier and write to the appropriate tier. As data ages, the software moves the data to a more appropriate tier. The storage tiering occurs without the application knowing it is occurring, making the tiering transparent to the application.

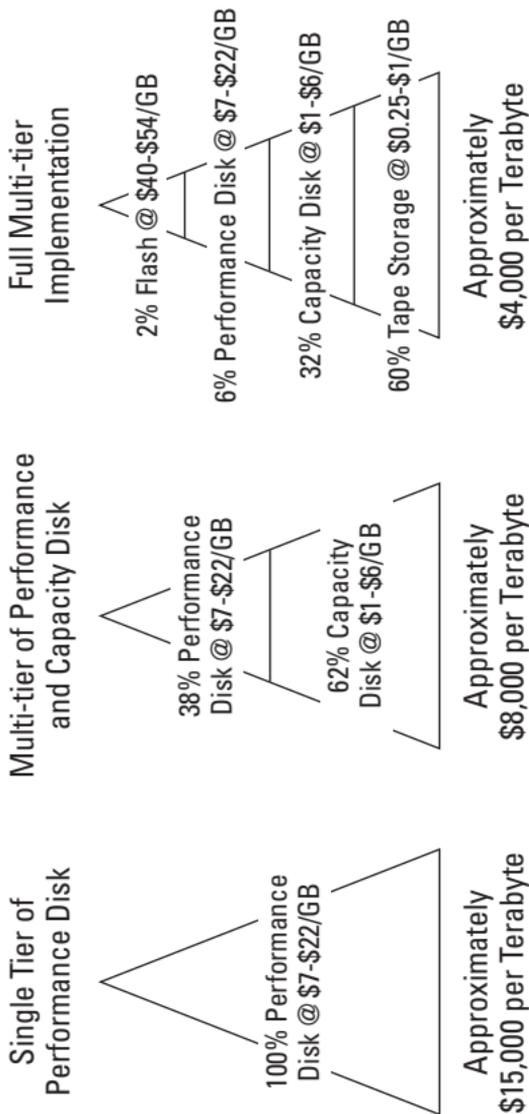
## *Balancing Storage Cost, Speed, and Capacity*

So, how does storage tiering save money over utilizing a single storage tier for all data? The key is that by placing all less-used, non-critical data on lower-cost storage tiers, it is not necessary to store archive data (which you may never even use) on high-performance (and higher-priced) storage media.

Performance is improved because the frequently accessed data is placed on the high-performance flash or performance disk tier, and the underlying configuration of the flash and performance disk tiers can be defined to maximize the performance of that data. Although high-performance flash or disk is more expensive, because only a small percentage of the total storage is high performance, the total storage cost is still reduced.

Total capacity needs to grow as total storage requirements increase at an explosive rate. Fortunately, most growth is within the capacity disk and long-term storage tiers because most of the data growing is seldom, if ever, used. Storing this data on less-expensive capacity disks and archive tiers keeps total storage costs manageable.

In Figure 2-1, you see a comparison of how an intelligently tiered solution can drive down storage costs.



**Figure 2-1:** Tiered-storage cost savings.

Source: *Horison Information Strategies 2010*

In the first scenario, a single tier of a primary storage-performance disk was used for the entire organization. While this would provide fast access for the data, much of that data would seldom or never be used, so access speed is unappreciated. Average cost is approximately \$15,000 per terabyte.

The second scenario represents a very common implementation. A mix of a 38 percent high primary storage performance disk is used with a 62 percent secondary storage capacity disk. Using RAID technologies to create high-performance and capacity-storage disk tiers is common and has driven down the average cost to approximately \$8,000 per terabyte.

The third scenario shows a properly tiered implementation using flash storage, a high primary storage performance disk, a secondary storage capacity disk, and long-term tape storage. This reduced the average cost per terabyte to \$4,000.

With intelligent storage tiering, storage costs were reduced by nearly 75 percent from an average of \$15,000 per terabyte to \$4,000 per terabyte. This was done while still meeting performance and capacity requirements — which is an impressive feat!



## Chapter 3

---

# Oracle's Complete Tiered Storage Solutions

.....

### *In This Chapter*

- ▶ Covering the spectrum of storage needs
  - ▶ Implementing storage for each tier
  - ▶ Putting together multiple storage tiers
- .....

Near limitless options exist for implementing tiered storage. Oracle provides a large range of hardware and software tiering technologies to realize the goal of tiered storage: performance and capacity at a lower cost. Beyond simply providing storage devices, Oracle provides valuable storage integration components, dedicated file-system tiering devices, and application-level tiering products to provide a complete storage tiering solution.

In this chapter, we examine the wide range of Oracle products available to implement tiered storage, their technologies, and how they can be implemented in real-world situations.

## ***Complete Coverage of Tiered Storage Needs***

With Oracle's acquisition of Sun Microsystems, Oracle now commands complete coverage of hardware and software tiering solutions. Originally focused on software and database-centric tiering products, Oracle already possessed extremely capable tiering products. Sun brought to the table its highly capable hardware and file-system-level storage tiering components. Combined, the Oracle and Sun storage tiering capabilities offer the best-in-breed database, software, file-system, and hardware storage tiering components.

Furthermore, now that Oracle owns both the hardware and software, those previously separate components are engineered to work with each other to gain higher-performance optimization. Oracle realizes that integration and optimization of complete solutions offer far larger benefits than highly optimized individual components. Tighter integration of the tiered storage components allows for greater efficiency, flexibility, and ultimately lower total cost.

## ***Correct Storage for Each Level of Tiering***

We identify storage technologies as flash storage, performance disk, capacity disk, and tape. We also identify application-transparent tiering as file-system-level tiering without knowledge of the application. Application-aware tiering products use additional

knowledge of the data, formats, and usage patterns to explicitly manage storage tiers. Oracle provides both application-transparent and application-aware tiering for flash, disk, and tape.

## ***Oracle's Sun Flash Storage***

At the top tier for storage performance, Oracle's Sun Flash Storage is the fastest in the industry. Using Sun FlashFire technology, Oracle's Sun Storage F5100 Flash Array is more than three orders of magnitude faster than disk, more compact than disk, and uses less power than disk.

Designed to eliminate I/O bottlenecks, Sun FlashFire technology doesn't have moving parts, which allows it to run much faster, cooler, and be smaller than even the fastest disk arrays. An additional benefit of not having moving parts, flash storage solutions are less prone to mechanical failure than disks.

The Sun F5100 Flash Array offers up to 1.6 million I/O Operations Per Second (IOPS) and up to 1.9 terabytes of storage in a single 1.75-inch-high chassis. The IOPS performance of a single Sun Storage F5100 Flash Array is equivalent to that of more than 3,000 hard disk drives, and performance and capacity in a solution can be increased by simply adding more flash arrays.

The Sun Flash Accelerator F20 PCIe card is a flash card that provides vast performance over disk at a fraction of the size and is easy to implement.

Because Sun Flash Storage technology is much faster than disk, Oracle has engineered its database technology to leverage this speed, as we discuss in later sections of this chapter.

## ***Oracle's Sun 6000 Storage Arrays***

Sun 6000 Storage Arrays are performance disk storage arrays for your highly accessed, mission-critical data. The Oracle Sun 6000 Storage Array family has models sized for small, medium and large installations.

Built on a reliable and scalable architecture, many organizations use the Sun 6000 Storage Arrays as their primary tier of storage. The 6000 family offers advanced data services such as snapshots, clones, replication, and encryption. Furthermore, flash and capacity disk capabilities can also be integrated directly into these arrays. The arrays are connected to individual servers or a SAN through low-latency, high-performance Fibre Channel. Supported storage drives can be Fibre Channel (FC), Encrypted Fibre Channel, Serial Advanced Technology Attachment (SATA), or Solid State Disk (SSD); and you can intermix these drives within the same disk tray. Common RAID configurations are supported. Management is browser-based via Sun Storage Common Array Manager (CAM) software.

## ***Oracle's Sun ZFS Storage Appliance***

Oracle's Sun ZFS Storage Appliances are unified storage solutions that both support file serving and block access protocols to simplify deployment and minimize total cost of ownership (TCO). These solutions offer high throughput performance and are often used to meet capacity storage requirements in a tiered storage environment, or to utilize their own internal storage tiering to support complete storage installations.

A key feature of these devices is the use of Hybrid Storage Pools (HSP). HSPs are a form of internal storage tiering used within Sun ZFS Storage Appliances to

integrate memory, flash drive technology, and capacity disk. Based on usage patterns, the Oracle Solaris ZFS file system automatically shifts I/O for specific files to faster or slower storage media components with the HSP to maintain a fast overall response time. By maintaining a pool of different storage media, Oracle Solaris ZFS is able to move data to the high-performance media to provide performance while using less expensive media in the background.

Additional optimizations to reduce the amount of data physically stored include in-line deduplication and compression.

Management of these devices via a Web-browser interface is simplified to reduce administrative costs and accelerate implementation timeframes. DTrace Analytics is a powerful real-time software package that provides graphical analysis and monitoring capabilities. These features assist the administrator to better manage performance and prevent issues.

## ***Oracle's StorageTek Tape***

Tape isn't just for backups of disk; tape can be used within a tiered storage environment to provide long-term and energy efficient storage. The combination of Oracle's Sun Storage Archive Manager software with Oracle's StorageTek tape libraries allows data to be accessed from standard, disk-oriented applications, despite the fact that it's not on disk. Because tape is reliable, uses less energy than disk, and has a far cheaper price per GB than disk, tape is the foundation of the long-term storage tier as well as for data archives.

Oracle offers the time-tested StorageTek suite of tape libraries, tape drives, tape media, and tape

management software. Scalable and versatile, the high performance StorageTek products offer advanced consolidation and media reuse capabilities to lower cost. StorageTek Virtual Storage Manager (VSM) allows mainframe computers to take advantage of tiered storage capabilities by transparently utilizing two tiers of disk and one tier of tape to provide scalability and performance.

Security with tape is enhanced with Oracle Key Manager (OKM), which provides high-speed encryption of your tape drives and dedicated encryption key management. Compliant with U.S. Federal Information Processing Standard (FIPS) 140-2 Level 3 standard, this can meet regulatory requirements for encrypted data at rest.

## ***Integrating Multiple Storage Tiers***

The greatest benefits of storage tiering are available to organizations that integrate the tiers to provide efficient data management that span multiple storage technologies. In this section, we look at Oracle's integrated tiered storage offerings.

### ***Oracle's Sun Storage Archive Manager (SAM)***

Oracle's Sun Storage Archive Manager (SAM) classifies and manages data located on different storage tiers based on data management policies. SAM will archive inactive data from high-performance storage tiers to less active tiers such as capacity disk or tape. Predefined policies like high-water marks determine

what data is or isn't archived, to what tier of storage, and when.

SAM provides the automated capability to manage the placement of data on the appropriate storage tier in an application transparent manner.

SAM can be incorporated with additional Oracle products such as Oracle Database Recovery Manager (RMAN) to assist managing database backups. An example would be moving backup images to archive tape storage even though RMAN believes it is operating on disk.

## ***Oracle Database 11g***

Every new version of Oracle Database offers new enhancements to performance, storage, and management. Some features are new; others are improvements on existing features. Oracle Database 11g offers several key features to be aware of that leverage flash storage and tiered storage.

### ***Smart Flash Cache***

Smart Flash Cache is an enhancement in Oracle Database 11g Release 2 to extend the database instance's buffer cache area outside the normal memory confines and onto flash storage.

The database buffer cache is a memory pool within Oracle Database instances. Key to performance is data being *cached*, or stored temporarily in a block of memory for easy access, in the database buffer to avoid costly disk I/O calls. Previously, the database buffer cache was a memory-only structure with limited size. With Oracle Database 11g Release 2, the buffer cache can be transparently extended into flash storage. This allows the buffer cache to be large enough to

support the needs of the database instance. With this improvement, performance is improved because rather than utilizing expensive disk I/O, data is cached on quick and efficient flash storage.

### *Index file optimizations*

Indexes are critical to fast access to data in large tables; indeed, without indexing, large databases would not be functional. Traditionally, indexes have existed on the same disk storage medium as data tables. One best practice for increasing database performance is to place index files on a tier of flash storage. The enhanced speed of flash over disk allows for increased index access and therefore faster data access.



Using Oracle tools such as Statspack and Automatic Workload Repository reports allows the DBA to identify high-demand data best served by flash storage.

### *Partitioning*

Data stored in database tables can reach extremely large sizes that slow down many database systems. Indexing key data elements to allow faster searches provides critical improvement, but does not go far enough. Database partitioning can be used to move massive amounts of older, less-used data away from newer, more active data and allows for a much smaller amount of data to be searched in a more optimal manner.

Partitioning enables Oracle database storage optimization by creating portable tables to place massive, older data partitions on less expensive tiers while the smaller, more active data set goes on high-performance tiers. This configuration provides faster

access for the most commonly used data and makes wise use of less expensive storage tiers.

An optimal configuration is a combination of flash storage, performance disk, and capacity disk, assigned to the appropriate partitions based on data age and usage.



Oracle Partitioning is not new; it has been around for several versions of Oracle Database. With each new major database release, existing database options are improved in terms of performance, capability, and manageability. Make the time to review the enhancements so you can take full advantage of them.

### ***Oracle Advanced Compression***

Oracle Database Advanced Compression is a technology to compress database data to reduce hardware storage requirements. In some cases, the reduction can be a factor of 2-4X. Compressing data in database tables provides reductions in storage across the production, test, and development environments. Furthermore, because the databases are smaller, the amount of space required for backup storage is reduced. These reductions can mean real savings in hardware storage.

### ***Oracle Exadata Database Machine***

Oracle Exadata Database Machine is Oracle's largest, most highly scalable and powerful database appliance. A self-contained database powerhouse, it is Oracle's flagship product for large, powerful online transaction processing (OLTP) and data warehousing.

Exadata is a database appliance consisting of servers, InfiniBand switches, flash storage, capacity disk, Oracle

developed and supported operating systems, and Oracle RAC (Real Application Cluster) database software. Exadata implements storage tiering within its storage nodes through the use of Oracle Exadata Smart Flash Cache. Exadata exemplifies the very tenets of tiered storage in a complete package that is massively scalable, secure, and redundant.

## Chapter 4

---

# Implementing Oracle Tiered Storage

.....

### *In This Chapter*

- ▶ Identifying who benefits the most
  - ▶ Looking at actual uses
  - ▶ Deciding on your product stack
  - ▶ Configuring, installing, and getting support
- .....

**T**aking tiered storage from theory to a real-world implementation is how the benefits of improved performance and lower cost are realized. It isn't enough to understand what storage *can* do; you need to see how it works in real-world implementations.

In this chapter, we identify environments that benefit most from Oracle tiered storage. Here, we discuss sample implementations to show how tiered storage is implemented in several different environments. We also identify what you need to know about product selection, implementation, and support of Oracle tiered storage.

## *Environments That Benefit the Most from Tiered Storage*

Given the scope of Oracle's tiered storage solutions and their integration with so many products, it would be challenging to find an environment that could *not* benefit from Oracle tiered storage.

Supporting both application-aware and application-transparent tiering components, it is easy to implement storage tiering into your environment. Application-aware products such as Oracle databases lend themselves to tiering very easily. File system-based tiering solutions cover virtually any environment because the applications aren't aware of the storage tiering that is occurring.



You don't have to be an "Oracle Database shop" to take advantage of Oracle tiered storage. File system-based tiering with Oracle Sun devices allows you to take advantage of Oracle storage tiering without even owning a database!

These environments fit easily into Oracle tiered storage:

- ✓ Environments storing unstructured data are growing at an explosive rate. Oracle tiered storage can store this data on less expensive storage devices, and Oracle Content Management can catalog and manage it.
- ✓ High performance Oracle database environments can leverage Oracle Flash Storage products. Flash storage is an extremely fast-and-efficient storage medium used to store your frequently used data.

- ✓ Large and growing Oracle Database instances. Oracle Database can use partitioning to store your seldom-accessed, large amounts of data on capacity disks and longer term tape archive storage. This action reduces your storage costs.
- ✓ Mainframe systems utilizing tape libraries. Oracle StorageTek Virtual Storage Manager (VSM) provides intelligent virtual tape library appliances that use easy-to-manage policies to implement two levels of disk and one level of tape behind the VSM controller.
- ✓ Small and mid-size companies can take advantage of the performance and capacity disk tiers provided by Oracle Sun 6000 Storage Arrays and Sun ZFS Storage Appliances. Scalable as storage requirements grow, Oracle Sun 6000 Storage Arrays provide performance disk. Sun ZFS Storage Appliances implement storage tiering within the appliance to provide high-performance and high-capacity while minimizing costs.

## ***Oracle Tiered Storage Use Cases***

Understanding the technical details of products supporting tiered storage is useful, but seeing how these technologies can be used in a practical implementation really helps in realizing their capabilities.

### ***Database Performance optimization***

Oracle Database has been the premier database for large, critical applications for many years because it is fast and offers advanced features not found in other

database products. Two features that work well are Oracle Smart Flash Cache and partitioning.

Extending database buffer cache to flash storage reduces the amount of memory needed to support a database instance, yet improves performance by reducing costly disk I/O operations. Not all data that needs to be cached can always fit in memory, but using flash storage as part of the buffer cache can greatly improve performance. Furthermore, this is easy to implement because the database automatically manages data within the buffer cache for you.

On the other hand, leveraging database partitioning with flash storage will allow fast access to the most frequently used partitions and lower cost storage for less frequently accessed data. Be sure to allocate the flash storage for your index files and active partitions. Recent but less active data is partitioned on the performance disk storage tier. Older, seldom used data partitions are stored on capacity disks or archive storage tiers.

## ***Enterprise Content Management***

Oracle Enterprise Content Management Suite is an Oracle product centered on managing unstructured data. Unstructured data is data beyond characters or numbers; it is images, video, diagrams, and documents. Business examples would be engineering documents, legal contracts, and medical images.

Unstructured data is more difficult to manage because it is unlike the data you are accustomed to storing in a traditional database table setting. Furthermore, it's often not stored in a centralized database, but rather across different applications and servers.

Oracle Enterprise Content Management is a content management and tiered storage solution designed to handle unstructured data. Central to the solution is Oracle's Sun Storage Archive Manager. This suite allows you to find, store, version, search, and display unstructured data. The large cost benefit is that it brings all the unstructured data into a managed control system.

Oracle tiered storage management tools (such as Oracle's Sun Storage Archive Manager) can be incorporated with the Oracle Enterprise Content Management Suite to automatically move infrequently used unstructured data to the appropriate capacity disk or tape archive tier.



Many forms of unstructured data are prime candidates to be stored on capacity disk or tape archive because they seldom change and do not have a high number of IOPS per GByte of storage. Unlike structured data in a database table that can be updated with a simple SQL statement, how often does a video clip, MRI image, or legal contract stored as a PDF get updated? The lack of updates makes this data a candidate for less expensive storage.

## *Mainframe*

Oracle StorageTek Virtual Storage Manager (VSM) is used in the mainframe world to provide high performance and scalable tape management solutions. Acting as a virtual storage manager, I/O to tape from mainframe systems is optimized because instead of going directly to tape, the I/O is managed by VSM. This allows for enhanced efficiencies in throughput, scalability

with automated migration of data between two levels of disk and tape. For mainframe systems, Oracle StorageTek VSM is a way to intelligently manage storage.

## *Selecting the Product Stack*

Selecting your product stack depends on your technical architecture and the nature of the data you store. Knowing the technical architecture of your system and where it is heading in the future makes these decisions much easier.

- ✔ Are you an Oracle Database shop? If so, look at application-aware features of the Oracle 11g Release 2 database such as partitioning, compression, and using a buffer cache with flash.
- ✔ If you are an Oracle Database shop, the database is already storage tiering application-aware, so consider implementing partitions with higher performance flash and lower cost unified storage as part of your existing installation. Consider also supplementing that with file system-based tiering solutions, particularly for unstructured data.
- ✔ Is your data unstructured and growing rapidly? If so, consider Oracle Enterprise Content Management software combined with Sun Storage Archive Manager software.
- ✔ Do you want to take advantage of high-performance flash and capacity disk but without worrying about implementations? If so, Oracle's Sun (pp) ZFS Storage Appliances can do much of the work for you.

Knowing the data you have, it's structure, and how often it's requested, update characteristics and retention

period is critical. This allows you to further identify which storage tiering products are right for you.

Once you have an idea of where Oracle tiered storage solutions may benefit you, then do a performance and cost analysis. Do the solution(s) provide you enough performance gain, reduction in administration costs, and reduction in storage costs to justify purchase and implementation? That is the key question. And it is a question that you must ask in terms of where your organization is today and in the future.

## *Installation and Configuration*

The technology behind Oracle tiered storage is complex and amazing, but installation and configuration needn't be difficult. Follow these tips to be successful.

- ✔ Understand the environment you're implementing the component in and what you want it to do.
- ✔ Read the installation and configuration documentation. It sounds simple, but it is the best thing you can do for yourself.
- ✔ Stay within supported installation and configuration matrices. Oracle tests and recommends specific configurations so stay within those if possible.
- ✔ If you have a problem, ask for help! Oracle Support can help you get your system where it needs to be when you run into issues.

The technology behind Oracle tiered storage is impressive, but with proper preparation, planning, and careful work it can be implemented without major difficulty.

## *Support and Operations*

The same tips that get you through installation and configuration will serve you well with support and operations. Understanding how you want to use the products, reading the documentation, and asking questions are keys to success. If appropriate, consider taking a training course to better understand the capabilities of the technology.

Additionally, Oracle provides many rich administrative tools for its products; understand and use them. Take advantage of the automation features to make your life easier and the system more efficient.

## Chapter 5

# Ten (Okay, Five) Things to Look For in a Tiered Storage Solution

---

### *In This Chapter*

- ▶ Considering the move to tiered storage
- 

**H**ere are five key items you need to know.

### *Are You a Good Fit?*

You are a good fit if you fall into these categories.

- ✔ You have explosively expanding data growth.
- ✔ You use unstructured data such as images, video clips, and documents.
- ✔ You use Oracle databases, regardless of size.
- ✔ You simply want to run a far more efficient and less expensive IT organization.

## *Do You Know Where to Start?*

Begin by implementing tiered storage with your existing storage assets. As new storage media is purchased, integrate it into your tiered storage strategy.

## *Are You Using All the Options?*

Use all the tiering options available to you! Some tiering solutions are limited to disk only. Oracle storage tiering is flash, performance disk, capacity disk, and tape archive. Be sure to use all the tiering mediums available to get the largest impact.

## *Are You Using Software Wisely?*

Let the software do the work for you! Designed to simplify administration, automate tasks, and improve availability and performance, these are the tools to use.

## *Are You Avoiding Common Barriers?*

Avoid these barriers to Oracle tiered storage.

- ✓ Don't believe that tiered storage is for someone else; it is for you and your organization!
- ✓ Don't try for a massive implementation at once; those seldom occur. Implement in milestones and components to ensure success.
- ✓ Do track performance gains and cost reductions!

# Achieving performance while reducing storage costs with tiered storage!

Explosive growth in storage requirements due to technology shifts, unstructured data, and regulatory mandates is more than traditional storage techniques can handle. IT budgets cannot support the growth necessary without a change in storage practices. Tiered storage is the smart solution for meeting the new storage challenges without skyrocketing costs.

- **Understand why current storage strategies are not sustainable for growing data pools** — find out why single tier solutions cannot support future growth and why tiered storage is far superior in performance and cost
- **Discover tiered storage architecture** — learn how multiple storage tiers work based on data access requirements
- **Move toward enterprise implementation** — see how storage tiering fits into your architecture right now using both application-aware and file system-based storage tiering solutions

**Oracle** (NASDAQ: ORCL) is the world's most complete, open, and integrated business software and hardware systems company. For more information about Oracle, visit [oracle.com](http://oracle.com).



## Open the book and find:

- What storage tiering means
- How tiered storage works and why it makes sense
- Ways tiered storage fits in your current system
- Why file-system tiering may be easier than you think

**Go to [Dummies.com](http://Dummies.com)**  
for videos, step-by-step examples,  
how-to articles, or to shop!

For Dummies®  
A Branded Imprint of



ISBN 978-1-118-06262-3  
Book not for resale