



DATABASE ADMINISTRATION

The Complete Guide to DBA Practices and Procedures

SECOND EDITION

CRAIG S. MULLINS

Database Administration

The Complete Guide to DBA Practices and Procedures

Second Edition

Craig S. Mullins

◆ Addison-Wesley

Upper Saddle River, NJ • Boston • Indianapolis • San Francisco
New York • Toronto • Montreal • London • Munich • Paris • Madrid
Capetown • Sydney • Tokyo • Singapore • Mexico City

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and the publisher was aware of a trademark claim, the designations have been printed with initial capital letters or in all capitals.

The author and publisher have taken care in the preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

The publisher offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please contact:

U.S. Corporate and Government Sales
(800) 382-3419
corpsales@pearsontechgroup.com

For sales outside the United States, please contact:

International Sales
international@pearsoned.com

Visit us on the Web: informit.com/aw

Library of Congress Cataloging-in-Publication Data

Database administration : the complete guide to DBA practices and procedures /

Craig S. Mullins.—2 [edition].

pages cm

Includes bibliographical references and index.

ISBN 978-0-321-82294-9 (alk. paper)—ISBN 0-321-82294-3 (alk. paper)

1. Database management. I. Title.

QA76.9.D3M838 2013

005.74—dc23

20120294

Copyright © 2013 Craig S. Mullins

All rights reserved. Printed in the United States of America. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. To obtain permission to use material from this work, please submit a written request to Pearson Education, Inc., Permissions Department, 200 Old Tappan Road, Old Tappan, New Jersey 07675, or you may fax your request to (201) 236-3290.

ISBN-13: 978-0-321-82294-9

ISBN-10: 0-321-82294-3

Text printed in the United States on recycled paper at RR Donnelley in Crawfordsville, Indiana.

Second printing, November 2015

Editor-in-Chief

Mark Taub

Acquisitions Editor

Greg Doench

Development Editor

Susan Brown Zahn

Managing Editor

John Fuller

Production Editor

Caroline Senay

Copy Editor

Barbara Wood

Indexer

Richard Evans

Proofreader

Diane Freed

Technical Reviewers

William Arledge

Kevin Kline

Editorial Assistant

Michelle Housley

Cover Designer

Chuti Prasertsith

Compositor

Rob Mauhar

The CIP Group

Accolades for *Database Administration*

“I’ve forgotten how many times I’ve recommended this book to people. It’s well written, to the point, and covers the topics that you need to know to become an effective DBA.”

—Scott Ambler, Thought Leader, Agile Data Method

“This is a well-written, well-organized guide to the practice of database administration. Unlike other books on general database theory or relational database theory, this book focuses more directly on the theory and reality of database administration as practiced by database professionals today, and does so without catering too much to any specific product implementation. As such, *Database Administration* is very well suited to anyone interested in surveying the job of a DBA or those in similar but more specific roles such as data modeler or database performance analyst.”

—Sal Ricciardi, Program Manager, Microsoft

“One of Craig’s hallmarks is his ability to write in a clear, easy-to-read fashion. The main purpose of any technical book is to transfer information from writer to reader, and Craig has done an excellent job. He wants the reader to learn—and it shows.”

—Chris Foot, Manager, Remote DBA Experts and Oracle ACE

“A complete and comprehensive listing of tasks and responsibilities for DBAs, ranging from creating the database environment to data warehouse administration, and everything in between.”

—Mike Tarrani, Computer Consultant

“I think every business manager and every IT manager should have a copy of this book.”

—Dan Hotka, Independent Consultant and Oracle ACE

“This book by Craig Mullins is wonderfully insightful and truly important. Mullins describes the role and duties of data administrators and database administrators in modern organizations with remarkable insight

and clarity.”

—Michael Tozer, Author and former U.S. Navy officer

To my wife, Beth, for her unending love, constant support, and beautiful smile.

Contents

Preface

[How to Use This Book](#)

Acknowledgments

About the Author

Chapter 1 What Is a DBA?

[Why Learn Database Administration?](#)

[A Unique Vantage Point](#)

[DBA Salaries](#)

[Database Technology](#)

[The Management Discipline of Database Administration](#)

[A Day in the Life of a DBA](#)

[Evaluating a DBA Job Offer](#)

[Database, Data, and System Administration](#)

[Data Administration](#)

[Database Administration](#)

[System Administration](#)

[DBA Tasks](#)

[Database Design](#)

[Performance Monitoring and Tuning](#)

[Ensuring Availability](#)

[Database Security and Authorization](#)

[Governance and Regulatory Compliance](#)

[Backup and Recovery](#)

[Ensuring Data Integrity](#)

[DBMS Release Migration](#)

[Jack-of-All-Trades](#)

The Types of DBAs

System DBA

Database Architect

Database Analyst

Data Modeler

Application DBA

Task-Oriented DBA

Performance Analyst

Data Warehouse Administrator

Staffing Considerations

How Many DBAs?

DBA Reporting Structures

Multiplatform DBA Issues

Production versus Test

The Impact of Newer Technology on DBA

Procedural DBAs: Managing Database Logic

The Internet: From DBA to eDBA

The Personal DBA and the Cloud

NoSQL, Big Data, and the DBA

New Technology Impacts on DBA

DBA Certification

The Rest of the Book

Review

Bonus Question

Chapter 2 Creating the Database Environment

Defining the Organization's DBMS Strategy

Choosing a DBMS

DBMS Architectures

DBMS Clustering

DBMS Proliferation

[Hardware Issues](#)

[Cloud Database Systems](#)

[Installing the DBMS](#)

[DBMS Installation Basics](#)

[Hardware Requirements](#)

[Storage Requirements](#)

[Memory Requirements](#)

[Configuring the DBMS](#)

[Connecting the DBMS to Supporting Infrastructure Software](#)

[Installation Verification](#)

[DBMS Environments](#)

[Upgrading DBMS Versions and Releases](#)

[Features and Complexity](#)

[Complexity of the DBMS Environment](#)

[Reputation of the DBMS Vendor](#)

[Support Policies of the DBMS](#)

[Organization Style](#)

[DBA Staff Skill Set](#)

[Platform Support](#)

[Supporting Software](#)

[Fallback Planning](#)

[Migration Verification](#)

[The DBMS Upgrade Strategy](#)

[Database Standards and Procedures](#)

[Database Naming Conventions](#)

[Other Database Standards and Procedures](#)

[DBMS Education](#)

[Summary](#)

[Review](#)

[Bonus Question](#)

[Suggested Reading](#)

[Chapter 3 Data Modeling and Normalization](#)

[Data Modeling Concepts](#)

[*Entity-Relationship Diagramming*](#)

[The Components of a Data Model](#)

[*Entities*](#)

[*Attributes*](#)

[*Keys*](#)

[*Relationships*](#)

[Discovering Entities, Attributes, and Relationships](#)

[Conceptual, Logical, and Physical Data Models](#)

[What Is Normalization?](#)

[The Normal Forms](#)

[*First Normal Form*](#)

[*Second Normal Form*](#)

[*Third Normal Form*](#)

[*A Normalized Data Model*](#)

[*Further Normal Forms*](#)

[Normalization in Practice](#)

[Additional Data Modeling Issues](#)

[Summary](#)

[Review](#)

[*Bonus Question*](#)

[Suggested Reading](#)

[Chapter 4 Database Design](#)

[From Logical Model to Physical Database](#)

[*Transform Entities to Tables*](#)

[*Transform Attributes to Columns*](#)

[*Build Referential Constraints for All Relationships*](#)

[*Build Physical Data Structures*](#)

Database Performance Design

Designing Indexes

Hashing

Clustering

Interleaving Data

Denormalization

When to Denormalize

Prejoined Tables

Report Tables

Mirror Tables

Split Tables

Combined Tables

Redundant Data

Repeating Groups

Derivable Data

Hierarchies

Special Physical Implementation Needs

Denormalization Summary

Views

Data Definition Language

Temporal Data Support

A Temporal Example

Business Time and System Time

Summary

Review

Bonus Question

Suggested Reading

Chapter 5 Application Design

Database Application Development and SQL

SQL

[*Set-at-a-Time Processing and Relational Closure*](#)

[*Embedding SQL in a Program*](#)

[*SQL Middleware and APIs*](#)

[*Application Infrastructure*](#)

[*Object Orientation and SQL*](#)

[*Types of SQL*](#)

[*SQL Coding for Performance*](#)

[*Querying XML Data*](#)

[**Defining Transactions**](#)

[*Transaction Guidelines*](#)

[*Unit of Work*](#)

[*Transaction Processing Systems*](#)

[*Application Servers*](#)

[**Locking**](#)

[*Types of Locks*](#)

[*Lock Time-outs*](#)

[*Deadlocks*](#)

[*Lock Duration*](#)

[*Lock Escalation*](#)

[*Programming Techniques to Minimize Locking Problems*](#)

[*Locking Summary*](#)

[**Batch Processing**](#)

[**Summary**](#)

[**Review**](#)

[*Bonus Question*](#)

[**Suggested Reading**](#)

Chapter 6 Design Reviews

[**What Is a Design Review?**](#)

[*Rules of Engagement*](#)

[*Design Review Participants*](#)

Knowledge and Skills Required

Types of Design Reviews

Conceptual Design Review

Logical Design Review

Physical Design Review

Organizational Design Review

SQL and Application Code Design Review

Pre-Implementation Design Review

Post-Implementation Design Review

Design Review Output

Additional Considerations

Dealing with Remote Staff

Mentorship and Knowledge Transfer

Summary

Review

Suggested Reading

Chapter 7 Database Change Management

Change Management Requirements

The Change Management Perspective of the DBA

Types of Changes

DBMS Software

Hardware Configuration

Logical and Physical Design

Applications

Physical Database Structures

Impact of Change on Database Structures

The Limitations of ALTER

Database Change Scenarios

Comparing Database Structures

Requesting Database Changes

[Standardized Change Requests](#)

[Communication](#)

[Coordinating Database and Application Changes](#)

[Compliance](#)

[DBA Scripts and Change Management](#)

[Summary](#)

[Review](#)

[Suggested Reading](#)

[Chapter 8 Data Availability](#)

[Defining Availability](#)

[Increased Availability Requirements](#)

[Cost of Downtime](#)

[How Much Availability Is Enough?](#)

[Availability Problems](#)

[Loss of the Data Center](#)

[Network Problems](#)

[Loss of the Server Hardware](#)

[Disk-Related Outages](#)

[Operating System Failure](#)

[DBMS Software Failure](#)

[Application Problems](#)

[Security and Authorization Problems](#)

[Corruption of Data](#)

[Loss of Database Objects](#)

[Loss of Data](#)

[Data Replication and Propagation Failures](#)

[Severe Performance Problems](#)

[Recovery Issues](#)

[DBA Mistakes](#)

[Outages: Planned and Unplanned](#)

Ensuring Availability

Perform Routine Maintenance While Systems Remain Operational

Automate DBA Functions

Exploit High-Availability Features

Exploit Clustering Technology

Database Architecture and NoSQL

Summary

Review

Suggested Reading

Chapter 9 Performance Management

Defining Performance

A Basic Database Performance Road Map

Monitoring versus Management

Reactive versus Proactive

Preproduction Performance Estimation

Historical Trending

Service-Level Management

Types of Performance Tuning

System Tuning

Database Tuning

Application Tuning

Performance Tuning Tools

DBMS Performance Basics

Summary

Review

Bonus Question

Suggested Reading

Chapter 10 System Performance

The Larger Environment

Interaction with the Operating System

[*Allied Agents*](#)

[*Hardware Configuration*](#)

[*Components of the DBMS*](#)

[DBMS Installation and Configuration Issues](#)

[*Types of Configuration*](#)

[*Memory Usage*](#)

[*Data Cache Details*](#)

[*“Open” Database Objects*](#)

[*Database Logs*](#)

[*Locking and Contention*](#)

[*The System Catalog*](#)

[*Other Configuration Options*](#)

[*General Advice*](#)

[System Monitoring](#)

[Summary](#)

[Review](#)

[*Bonus Question*](#)

[Suggested Reading](#)

Chapter 11 Database Performance

[Techniques for Optimizing Databases](#)

[*Partitioning*](#)

[*Raw Partition versus File System*](#)

[*Indexing*](#)

[*Denormalization*](#)

[*Clustering*](#)

[*Interleaving Data*](#)

[*Free Space*](#)

[*Compression*](#)

[*File Placement and Allocation*](#)

[*Page Size \(Block Size\)*](#)

[Database Reorganization](#)

[*Determining When to Reorganize*](#)

[*Automation*](#)

[Summary](#)

[Review](#)

[Suggested Reading](#)

[Chapter 12 Application Performance](#)

[Designing Applications for Relational Access](#)

[Relational Optimization](#)

[*CPU and I/O Costs*](#)

[*Database Statistics*](#)

[*Query Analysis*](#)

[*Joins*](#)

[*Access Path Choices*](#)

[Additional Optimization Considerations](#)

[*View Access*](#)

[*Query Rewrite*](#)

[*Rule-Based Optimization*](#)

[Reviewing Access Paths](#)

[*Forcing Access Paths*](#)

[SQL Coding and Tuning for Efficiency](#)

[*A Dozen SQL Rules of Thumb*](#)

[*Additional SQL Tuning Tips*](#)

[*Identifying Poorly Performing SQL*](#)

[Summary](#)

[Review](#)

[Suggested Reading](#)

[Chapter 13 Data Integrity](#)

[Types of Integrity](#)

[Database Structure Integrity](#)

[*Types of Structural Problems*](#)

[*Managing Structural Problems*](#)

[Semantic Data Integrity](#)

[*Entity Integrity*](#)

[*Unique Constraints*](#)

[*Data Types*](#)

[*Default Values*](#)

[*Check Constraints*](#)

[*Triggers*](#)

[*Referential Integrity*](#)

[Temporal Database Systems](#)

[Summary](#)

[Review](#)

[Suggested Reading](#)

[Chapter 14 Database Security](#)

[Data Breaches](#)

[Database Security Basics](#)

[*Database Users*](#)

[Granting and Revoking Authority](#)

[*Types of Privileges*](#)

[*Granting to PUBLIC*](#)

[*Revoking Privileges*](#)

[*Label-Based Access Control*](#)

[*Security Reporting*](#)

[Authorization Roles and Groups](#)

[*Roles*](#)

[*Groups*](#)

[Other Database Security Mechanisms](#)

[*Using Views for Security*](#)

[*Using Stored Procedures for Security*](#)

[Encryption](#)

[*Data at Rest Encryption*](#)

[*Data in Transit Encryption*](#)

[*Encryption Techniques*](#)

[SQL Injection](#)

[*SQL Injection Prevention*](#)

[Auditing](#)

[External Security](#)

[*Job Scheduling and Security*](#)

[*Non-DBMS DBA Security*](#)

[DBMS Fixpacks and Maintenance](#)

[Summary](#)

[Review](#)

[Suggested Reading](#)

[Chapter 15 Regulatory Compliance and Database Administration](#)

[A Collaborative Approach to Compliance](#)

[*Why Should DBAs Care about Compliance?*](#)

[Metadata Management, Data Quality, and Data Governance](#)

[*Metadata*](#)

[*Data Quality*](#)

[*Data Governance*](#)

[Database Auditing and Data Access Tracking](#)

[*Database Auditing Techniques*](#)

[*Privileged User Auditing*](#)

[Data Masking and Obfuscation](#)

[*Data Masking Techniques*](#)

[Database Archiving for Long-Term Data Retention](#)

[*The Life Cycle of Data*](#)

[*Database Archiving*](#)

[*Components of a Database Archiving Solution*](#)

[*The Impact of e-Discovery on DBA*](#)
[Closer Tracking of Traditional DBA Tasks](#)

[*Database Change Management*](#)

[*Database Backup and Recovery*](#)

[Summary](#)

[Review](#)

[Suggested Reading](#)

[Chapter 16 Database Backup and Recovery](#)

[The Importance of Backup and Recovery](#)

[Preparing for Problems](#)

[Backup](#)

[*Full versus Incremental Backups*](#)

[*Database Objects and Backups*](#)

[*DBMS Control*](#)

[*Concurrent Access Issues*](#)

[*Backup Consistency*](#)

[*Log Archiving and Backup*](#)

[*Determining Your Backup Schedule*](#)

[*DBMS Instance Backup*](#)

[*Designing the DBMS Environment for Recovery*](#)

[*Alternate Approaches to Database Backup*](#)

[*Document Your Backup Strategy*](#)

[*Database Object Definition Backups*](#)

[Recovery](#)

[*Determining Recovery Options*](#)

[*General Steps for Database Object Recovery*](#)

[*Types of Recovery*](#)

[*Index Recovery*](#)

[*Testing Your Recovery Plan*](#)

[*Recovering a Dropped Database Object*](#)

[*Recovering Broken Blocks and Pages*](#)

[*Populating Test Databases*](#)

[Alternatives to Backup and Recovery](#)

[*Standby Databases*](#)

[*Replication*](#)

[*Disk Mirroring*](#)

[Summary](#)

[Review](#)

[Suggested Reading](#)

[Chapter 17 Disaster Planning](#)

[The Need for Planning](#)

[*Risk and Recovery*](#)

[General Disaster Recovery Guidelines](#)

[*The Remote Site*](#)

[*The Written Plan*](#)

[*Personnel*](#)

[Backing Up the Database for Disaster Recovery](#)

[*Tape Backups*](#)

[*Storage Management Backups*](#)

[*Other Approaches*](#)

[*Some Guidelines*](#)

[Disaster Prevention](#)

[*Disaster and Contingency Planning Web Sites*](#)

[Summary](#)

[Review](#)

[Suggested Reading](#)

[Chapter 18 Data and Storage Management](#)

[Storage Management Basics](#)

[Files and Data Sets](#)

[*File Placement on Disk*](#)

[*Raw Partitions versus File Systems*](#)

[*Temporary Database Files*](#)

[Space Management](#)

[*Data Page Layouts*](#)

[*Index Page Layouts*](#)

[*Transaction Logs*](#)

[Fragmentation and Storage](#)

[Storage Options](#)

[*RAID*](#)

[*JBOD*](#)

[*Storage Area Networks*](#)

[*Network-Attached Storage*](#)

[*Tiered Storage*](#)

[Planning for the Future](#)

[*Capacity Planning*](#)

[Summary](#)

[Review](#)

[Suggested Reading](#)

Chapter 19 Data Movement and Distribution

[Loading and Unloading Data](#)

[*The LOAD Utility*](#)

[*The UNLOAD Utility*](#)

[*Maintaining Application Test Beds*](#)

[EXPORT and IMPORT](#)

[Bulk Data Movement](#)

[*ETL Software*](#)

[*Replication and Propagation*](#)

[*Messaging Software*](#)

[*Other Methods*](#)

[Distributed Databases](#)

[*Setting Up a Distributed Environment*](#)

[*Data Distribution Standards*](#)

[*Accessing Distributed Data*](#)

[*Two-Phase COMMIT*](#)

[*Distributed Performance Problems*](#)

[Summary](#)

[Review](#)

[*Bonus Question*](#)

[Suggested Reading](#)

Chapter 20 Data Warehouse Administration

[What Is a Data Warehouse?](#)

[*Analytical versus Transaction Processing*](#)

[Administering the Data Warehouse](#)

[*Too Much Focus on Technology?*](#)

[*Data Warehouse Design*](#)

[*Data Movement*](#)

[*Data Cleansing*](#)

[*Data Warehouse Scalability*](#)

[*Data Warehouse Performance*](#)

[*Data Freshness*](#)

[*Data Content*](#)

[*Data Usage*](#)

[*Financial Chargeback*](#)

[*Backup and Recovery*](#)

[*Don't Operate in a Vacuum!*](#)

[Summary](#)

[Review](#)

[Suggested Reading](#)

Chapter 21 Database Connectivity

[Multitier, Distributed Computing](#)

[*A Historical Look*](#)

[*Business Issues*](#)

[*What Is Client/Server Computing?*](#)

[*Types of Client/Server Applications*](#)

[Network Traffic](#)

[*Database Gateways*](#)

[*Database Drivers*](#)

[*Connection Pooling*](#)

[Databases, the Internet, and the Web](#)

[*Internet-Connected Databases*](#)

[*Web Development and Web Services*](#)

[Summary](#)

[Review](#)

[Suggested Reading](#)

Chapter 22 Metadata Management

[What Is Metadata?](#)

[*From Data to Knowledge and Beyond*](#)

[*Metadata Strategy*](#)

[*Data Warehousing and Metadata*](#)

[Types of Metadata](#)

[Repositories and Data Dictionaries](#)

[*Repository Benefits*](#)

[*Repository Challenges*](#)

[*Data Dictionaries*](#)

[Summary](#)

[Review](#)

[Suggested Reading](#)

Chapter 23 DBA Tools

[Types and Benefits of DBA Tools](#)

[*Data Modeling and Design*](#)

Database Change Management

Table Editors

Performance Management

Backup and Recovery

Database Utilities

Data Protection, Governance, Risk, and Compliance Tools

Data Warehousing, Analytics, and Business Intelligence

Programming and Development Tools

Miscellaneous Tools

Examine Native DBA Tools

Evaluating DBA Tool Vendors

Homegrown DBA Tools

Summary

Review

Chapter 24 DBA Rules of Thumb

Write Down Everything

Keep Everything

Automate!

Share Your Knowledge

Analyze, Simplify, and Focus

Don't Panic!

Measure Twice, Cut Once

Understand the Business, Not Just the Technology

Don't Become a Hermit

Use All of the Resources at Your Disposal

Keep Up-to-Date

Invest in Yourself

Summary

Final Exam

Appendix A Database Fundamentals

[What Is a Database?](#)

[Why Use a DBMS?](#)

[*Advantages of Using a DBMS*](#)

[Summary](#)

[Appendix B The DBMS Vendors](#)

[The Big Three](#)

[The Second Tier](#)

[Other Significant Players](#)

[Open-Source DBMS Offerings](#)

[Nonrelational DBMS Vendors](#)

[NoSQL DBMS Vendors](#)

[Object-Oriented DBMS Vendors](#)

[PC-Based DBMS Vendors](#)

[Appendix C DBA Tool Vendors](#)

[The Major Vendors](#)

[Other DBA Tool Vendors](#)

[Data Modeling Tool Vendors](#)

[Repository Vendors](#)

[Data Movement and Business Intelligence Vendors](#)

[Appendix D DBA Web Resources](#)

[Usenet Newsgroups](#)

[Mailing Lists](#)

[Web Sites, Blogs, and Portals](#)

[*Vendor Web Sites*](#)

[*Magazine Web Sites*](#)

[*Consultant Web Sites*](#)

[*Blogs*](#)

[*Database Portals*](#)

[*Other Web Sites*](#)

Appendix E Sample DBA Job Posting

Job Posting

Database Administrator (DBA)

Bibliography

Database Management and Database Systems

Data Administration, Data Modeling, and Database Design

Database Security, Protection, and Compliance

Data Warehousing

SQL

Object Orientation and Database Management

Operating Systems

Related Topics

DB2

IMS

MySQL

Oracle

SQL Server

Sybase

Other Database Systems

Glossary

Index

Preface

The need for database administration is as strong as, or stronger than, it was when I originally wrote the first edition of this book in 2002. Relational database management systems are still at the core of most serious production systems, and they still need to be managed. And this is still the job of database administrators. Whether you use Oracle, Microsoft SQL Server, DB2, Informix, Sybase, MySQL, Teradata, PostgreSQL, Ingres, or any combination of these popular DBMS products, you will benefit from the information in this book.

But a decade is forever in the world of information technology. And even though some basic things stay the same (e.g., databases require administration), many things change. The second edition of this book incorporates the many changes that impact database administration that have occurred in the industry over the past decade. What made the book unique remains. It is still the industry's only non-product-based description of database administration techniques and practices. The book defines the job of database administrator and outlines what is required of a database administrator, or DBA, in clear, easy-to-understand language. The book can be used

- As a text for learning the discipline of database administration
- As the basis for setting up a DBA group
- To augment a DBMS-specific manual or textbook
- To help explain to upper-level management what a DBA is, and why the position is required

But what is new? One of the significant improvements added to this edition is coverage of regulatory compliance. The number of governmental and industry regulations has exploded over the course of the past decade, and many of these regulations dictate changes in the way that data is managed, handled, and processed. Although the most visible governmental regulation is undoubtedly the Sarbanes-Oxley Act (aka the U.S. Public Company Accounting Reform and Investor Protection Act of 2002), there are many others, including HIPAA (the Health Insurance Portability and Accountability Act) and GLB (the Gramm-Leach-Bliley Act) to name a

couple. The most visible industry regulation is PCI DSS (Payment Card Industry Data Security Standard). All of these regulations, and many others, impose an additional administrative burden on data. This edition of the book provides an entire chapter devoted to this topic, including the impact of regulatory compliance on data management tasks such as metadata management, data quality, database auditing, data masking, database archiving, and more traditional DBA tasks such as database change management and database recovery.

Database security is another rapidly evolving area that required a significant upgrade from the first edition. Fresh coverage is offered on new security functionality and requirements, including label-based access control, encryption, and preventing SQL injection attacks.

The book adds coverage of technology that was not widely adopted ten years ago, such as XML, and where appropriate it discusses nascent technology that DBAs should be aware of, including NoSQL and cloud computing. It also covers newer DBMS functionality, such as temporal database support and INSTEAD-OF triggers.

Finally, the entire book was reviewed and revised to ensure that each topic addressed up-to-date technology and requirements. Care was taken to ensure that the example DBMS features used to highlight specific technologies are accurate and up-to-date. For example, consider the descriptions of DB2 HADR, SQL Server 2012 AlwaysOn, and Oracle Transparent Data Encryption.

With the second edition of this book you now have a timely, accurate, and updated guide to implementing and maintaining heterogeneous database administration. You can use it to learn what is required to be a successful database administrator. And you can use it on the job in conjunction with the vendors' manuals or product-specific books for your particular DBMS products.

How to Use This Book

This book can be used as both a tutorial and a reference. The book is organized to proceed chronologically through DBA tasks that are likely to be encountered. Therefore, if you read the book sequentially from [Chapter 1](#) through [Chapter 24](#), you will get a comprehensive sequential overview of the DBA job. Alternatively, you can read any chapter independently because

each chapter deals with a single topic. References to other chapters are clearly made if other material in the book would aid the reader's understanding.

Acknowledgments

Writing is a rewarding task, but it also requires a lot of time—researching, writing, reviewing, editing, and rewriting over and over again until you get it just right. But no one can write a technical book in a vacuum. I had many knowledgeable and helpful people to assist me along the way.

First of all, I'd like to thank the many industry experts who reviewed the original book proposal. The following folks provided many useful suggestions and thoughts on my original outline that helped me to create a much better book: Michael Blaha, Keith W. Hare, Michael J. Hernandez, Robert S. Seiner, and David L. Wells. Additionally, I'd like to thank everyone who took the time to listen to my ideas for this book before I began writing. This list of folks is too numerous to include, and I'm sure I'd miss someone—but you know who you are.

I would like to thank the many folks who have reviewed and commented on the text of this book. For the second edition of the book, Bill Arledge and Kevin Kline provided their expertise to the review process and offered many helpful corrections and suggestions that improved the quality of the book. And let's not forget the reviewers of the first edition: Dan Hotka, Chris Foot, Chuck Kosin, David L. Wells, and Anne Marie Smith pored over each chapter of various incarnations of the manuscript, and this book is much better thanks to their expert contributions. Special thanks go to data modeling and administration gurus William J. Lewis and Robert S. Seiner, who took extra time to review and make suggestions on [Chapter 3](#). I'd also like to thank my brother, Scott Mullins, who offered his guidance on application design and development by reviewing [Chapter 5](#).

My appreciation goes to Mary Barnard, who did a wonderful job editing the first edition of this book; and Greg Doench, who did a similarly fantastic job with the second edition. Kudos to both Mary and Greg for making my book much more readable.

Additionally, thanks to the many understanding and patient folks at Addison-Wesley who worked with me to make each edition of the book come to fruition. This list includes Michelle Housley, Patrick Peterson, Stacie Parillo, Barbara Wood, and Mary O'Brien who were particularly helpful

throughout the process of coordinating the production of the book.

Thank you, too, to my wonderful wife, Beth, whose understanding and support made it possible for me to write this book. Indeed, thanks go out to all my family and friends for being supportive and helpful along the way.

And finally, a thank-you to all of the people with whom I have worked professionally at SoftwareOnZ, NEON Enterprise Software, Embarcadero Technologies, BMC Software, Gartner Group, PLATINUM Technology, Inc., Duquesne Light Company, Mellon Bank, USX Corporation, and ASSET, Inc. This book is a better one due to the many outstanding individuals with whom I have had the honor to work.

About the Author

Craig S. Mullins is President and Principal Consultant for Mullins Consulting, Inc., a consulting practice specializing in data management and database management systems. Craig has extensive experience in the field of database management, having worked as an application developer, a DBA, and an instructor with multiple database management systems, including DB2, Oracle, and SQL Server. Craig has worked in multiple industries, including manufacturing, banking, commercial software development, education, research, utilities, and consulting. Additionally, Craig worked as a Research Director with Gartner Group, covering the field of database administration. He is the author of *DB2 Developer's Guide*, the industry-leading book on DB2 for z/OS, currently in its sixth edition.

Craig is a frequent contributor to computer industry publications, having authored hundreds of articles in the past several years. His articles have appeared in popular industry magazines and Web sites, including *Database Programming & Design*, *Data Management Review*, *DBMS*, *DB2 Update*, *Oracle Update*, *SQL Server Update*, and many others. Craig writes several regular columns, including a monthly column called "The DBA Corner" for *Database Trends and Applications* magazine, a quarterly column called "The Database Report" for *The Data Administration Newsletter* (www.tdan.com), and a regular column on DB2 and mainframe data management called "z/Data Perspectives" for *zJournal Magazine*. Craig is also a regular blogger, managing and authoring two popular data-related blogs: *The DB2 Portal* (<http://db2portal.blogspot.com>) focusing on DB2 for z/OS and mainframe "stuff," and *Data and Technology Today* (<http://datatechnologytoday.wordpress.com>), which focuses on data and database management issues, DBA news and thoughts, metadata management, and data architecture, as well as data-related topics in the realm of IT and software. Craig is also the publisher and editor of *The Database Site* (www.thedatabasesite.com).

Craig regularly presents technical topics at database industry conferences and events. He has spoken to thousands of technicians about database management and administration issues at such conferences as Database and Client/Server World, SHARE, GUIDE, DAMA Symposium, Enterprise Data

World, IBM Information On Demand Conference, the DB2 Technical Conference, the International DB2 Users Group (IDUG), and Oracle Open World. He has also spoken at regional database user groups across North America, Europe, Asia, and Australia.

Craig graduated cum laude from the University of Pittsburgh with a double major in computer science and economics and a minor in mathematics. Craig has been appointed as an Information Management Champion by IBM for his work in the field of DB2 database administration, development, and management.

Readers can obtain information about this book, including corrections, future editions, and additional writings on database administration by the author, at the author's Web site at www.craigsmullins.com. The author can be contacted at craig@craigsmullins.com or in care of the publisher.

1. What Is a DBA?

Every organization that manages data using a database management system (DBMS) requires a database administration (DBA) group to ensure the effective use and deployment of the company's databases. And since most modern organizations of every size use at least one DBMS, the need for database administrators (DBAs) is greater today than ever before. However, the discipline of database administration is not well understood or universally practiced in a coherent and easily replicated manner.

The need for a database administrator is greater today than ever before.

There is a frequently repeated joke about database administration that helps to underscore both the necessity for DBA and lack of understanding of the DBA function. It goes something like this:

The CIO of Acme Corporation hires a management consulting company to help them streamline their IT operations. The consultant, determined to understand the way Acme works, begins by interviewing the CIO. One of the questions he asks is "So, I see that you have a DBA on staff; what does he do?"

The CIO says "Well, we use Oracle and I'm told that we need the DBA to make sure our Oracle databases stay online. I know that some of our critical business processes like order entry and inventory use Oracle, but I really don't know what the DBA does," says the CIO. "But please don't tell me I need another one because we can barely afford to pay the one we have!"

This is a sad, but too often true, commentary on the state of database administration in many organizations. Frequently the DBA is viewed as a guru or magician who uses tricks to make databases and systems operate efficiently. DBMS software is so complex these days that very few people understand more than just the basics (like SQL). But DBAs understand the complexities of the DBMS, making them a valuable resource. Indeed, sometimes the only source of database management and development

knowledge within the organization is the DBA.

The role of the DBA is as the guardian of the data as a corporate asset. So DBAs, in trying to protect the data, often are perceived as slow moving and risk adverse. Developers, on the other hand, are charged with building new applications and are constantly being challenged to build things fast and move on to the next project. Obviously, the difference between these two roles and expectations—with DBAs saying “change control, change management” and developers saying “deploy it now, deploy it now”—can create friction.

Another frequent criticism of the DBA staff is that they can be difficult to deal with. Sometimes viewed as prima donnas, DBAs can be curmudgeons who have vast technical knowledge but limited people skills. Just about every database programmer has a favorite DBA story. You know, those famous anecdotes that begin with “I have a problem . . .” and end with “. . . and then he told me to stop bothering him and read the manual.” DBAs simply do not have a “warm and fuzzy” image. This probably has more to do with the nature and scope of the job than anything else. The DBMS spans the enterprise, effectively placing the DBA on call for the applications of the entire organization.

The fact that DBAs often must sit down and work things through on their own can be a mitigating factor in this poor reputation. Many database problems require periods of quiet reflection and analysis to resolve. So DBAs do not generally like to be disturbed. But even though many problems will require solitude, there are many other problems that require a whole team to resolve. And due to the vast knowledge most DBAs possess, their quiet time is usually less than quiet; constant interruptions to answer questions and solve problems are a daily fact of life.

DBAs should not be encouraged to be antisocial. In fact, DBAs should be trained to acquire exceptional communication skills. Data is the lifeblood of computerized applications. Application programs are developed to read and write data, analyze data, move data, perform calculations using data, modify data, and so on. Without data there would be nothing for the programs to do. The DBA is at the center of the development life cycle—ensuring that application programs have efficient, accurate access to the corporation’s data. As such, DBAs frequently interface with many different types of people: technicians, programmers, end users, customers, and executives. However,

many DBAs are so caught up in the minutiae of the inner workings of the DBMS that they never develop the skills required to relate appropriately with their coworkers and customers.

DBAs need to acquire exceptional communication skills.

But we have not yet answered the question that is the title of this chapter: What is a DBA? The short answer to that question is simple: A DBA is the information technician responsible for ensuring the ongoing operational functionality and efficiency of an organization's databases and the applications that access those databases.

A DBA is the information technician responsible for ensuring the ongoing operational functionality and efficiency of an organization's databases and applications.

The long answer to that question requires a book to answer—this book. This text will define the management discipline of database administration and provide practical guidelines for the proper implementation of the DBA function. In order to begin to answer the question, though, this chapter will provide an introductory overview of database administration, covering the reasons to pursue a career as a DBA, high-level summaries of DBA tasks, and a look at the organizational structure of the DBA team.

Why Learn Database Administration?

As we have already mentioned, data is at the center of today's applications, and today's modern organization simply cannot operate without data. In many ways, business today *is* data.¹ Without data, businesses would not have the ability to manage finances, to conduct transactions, or to contact their customers. Databases are created to store and organize this data. The better the design and utility of the database, the better the organization will be positioned to compete for business.

Indeed, one of the largest problems faced by IT organizations is ensuring quality database administration. According to a recent study:²

- Good DBAs are hard to find and costly to acquire—76 percent of respondents say it takes more than three months to hire a DBA, and training a new DBA in a new environment takes months beyond that.

- The database infrastructure supporting applications is complex, chronically fragmented, and cumbersome to manage.

Both of these findings clearly indicate that database administration is a difficult job that is in high demand. Additionally, according to the Dice³ 2010–11 Tech Salary Survey, Oracle experience is requested in more than 15,000 job postings on any given day. Demand for Oracle skills is up 57 percent year over year, and the national average salary for technology professionals with experience in Oracle Database is \$90,914.

A Unique Vantage Point

The DBA is responsible for designing and maintaining an enterprise's databases, placing the DBA squarely at the center of the business. The DBA has the opportunity to learn about many facets of business and how they interrelate. The DBA can explore groundbreaking technologies as they are adopted by the organization. Exposure to new technology keeps the job stimulating—but frustrating if you are trying to figure out how a new technology works for the first time. The DBA is often working alone in these endeavors; he does not have access to additional expertise to assist when troubles arise. Therefore, a good DBA needs to enjoy challenges and be a good problem solver.

A good DBA needs to enjoy challenges and be a good problem solver.

DBA Salaries

As a technician you can find no more challenging job in IT than DBA. Fortunately, the job of DBA also is quite rewarding. DBAs are well paid. According to a salary study conducted by Global Knowledge⁴ and Tech-Republic,⁵ the average DBA salary is \$78,468, while their managers average \$87,261. These average salaries are a little lower than the numbers from the Janco⁶ salary survey (conducted the previous year).

For full-time employees functioning as DBAs, the mean salary ranges in the high \$80 thousands. Refer to [Table 1.1](#) for a breakdown of the mean salary for DBAs between 2006 and 2011 (according to the Dice 2010–11 Tech Salary Survey).

Table 1.1. Mean Salary for DBAs between 2006 and 2011

Year	Mean DBA Salary
2006-7	\$85,441
2007-8	\$85,092
2008-9	\$89,742
2009-10	\$91,283
2010-11	\$88,443

Keep in mind that the salary figures quoted here are for illustrative purposes only and will vary based on numerous factors. As might be expected, as the years of experience and number of people managed increases, so does the salary. Of course, DBA salaries, like all salaries, vary by region of the country as well. In the United States, DBA salaries are likely to be higher in the Northeast and on the West Coast than in other regions. Type of industry is also a factor, with pharmaceuticals paying higher than government entities (for example). When all of these factors are considered, the salary figures for DBAs are quite stable and on the higher end of IT individual contributors.

According to the 2012 Computerworld Salary Survey, the national average salary for a DBA is \$95,187, but the average salary for a DBA with 15 to 20 years of experience and working on the Pacific coast is \$103,597.⁷ Obviously experience level and location make a significant difference.

The U.S. Bureau of Labor Statistics (BLS)⁸ provides additional information about DBA employment and compensation. As of May 2010, the BLS reports that the median annual wages of database administrators were \$73,490 and the mean annual wage was \$75,730. The lowest 10 percent earned less than \$41,570, and the highest 10 percent earned more than \$115,660. The BLS also breaks down earnings statistics by geographical region, so you can use their data to determine expected salary ranges for your specific location of interest.

Perhaps more important than compensation is employability, and your prospects for employment as a DBA are quite good. According to the U.S. Bureau of Labor Statistics *Occupational Outlook Handbook*, 2010–11

Edition, “Employment is expected to grow much faster than the average, and job prospects should be excellent.” Indeed, between the years 2008 and 2018 the BLS estimates that the number of DBA jobs will increase by 20 percent.

The job market for DBAs will grow much faster than the average.

So, DBAs are well paid, highly employable, possess challenging jobs, and are likely to be engaged in the most visible and important projects. What’s not to like? Well, DBAs are expected to know everything, not just about database technology but about anything remotely connected to it. DBAs almost never work simple eight-hour days, instead working long days with a lot of overtime, especially when performance is suffering or development projects are behind schedule. According to industry analysts, the average DBA works more than 50 hours per week, including an average of six hours on weekends. DBAs frequently have to work on weekends and holidays to maintain databases during off-peak hours.

Database administration is a nonstop job.

The job of DBA is technically challenging and rewarding, but also potentially one that will exhaust and frustrate you. But don’t let that scare you. The positive aspects of the job far outweigh the negative.

Database Technology

The DBA is the IT professional who understands the specific details of database technology. It is important for the DBA to have a solid understanding of the fundamentals of database management. This requires much more than simple knowledge of rows, columns, and tables, or SQL.

This book assumes the reader has basic knowledge of relational database technology and DBMS fundamentals. For readers needing to review these concepts, please refer to [Appendix A](#), “[Database Fundamentals](#).” This is not a trivial matter. Sometimes the problem is that people think they know more than they actually do. For example, “[What is a database?](#)” I bet most people reading this believe they know the answer to that question. But some (perhaps many) of you would be wrong. SQL Server is not a database; it is a DBMS, or database management system. You can use SQL Server to create a database, but SQL Server, in and of itself, is not a database.

So what is a database? A database is an organized store of data wherein the data is accessible by named data elements (for example, fields, records, and files) (see [Figure 1.1](#)).

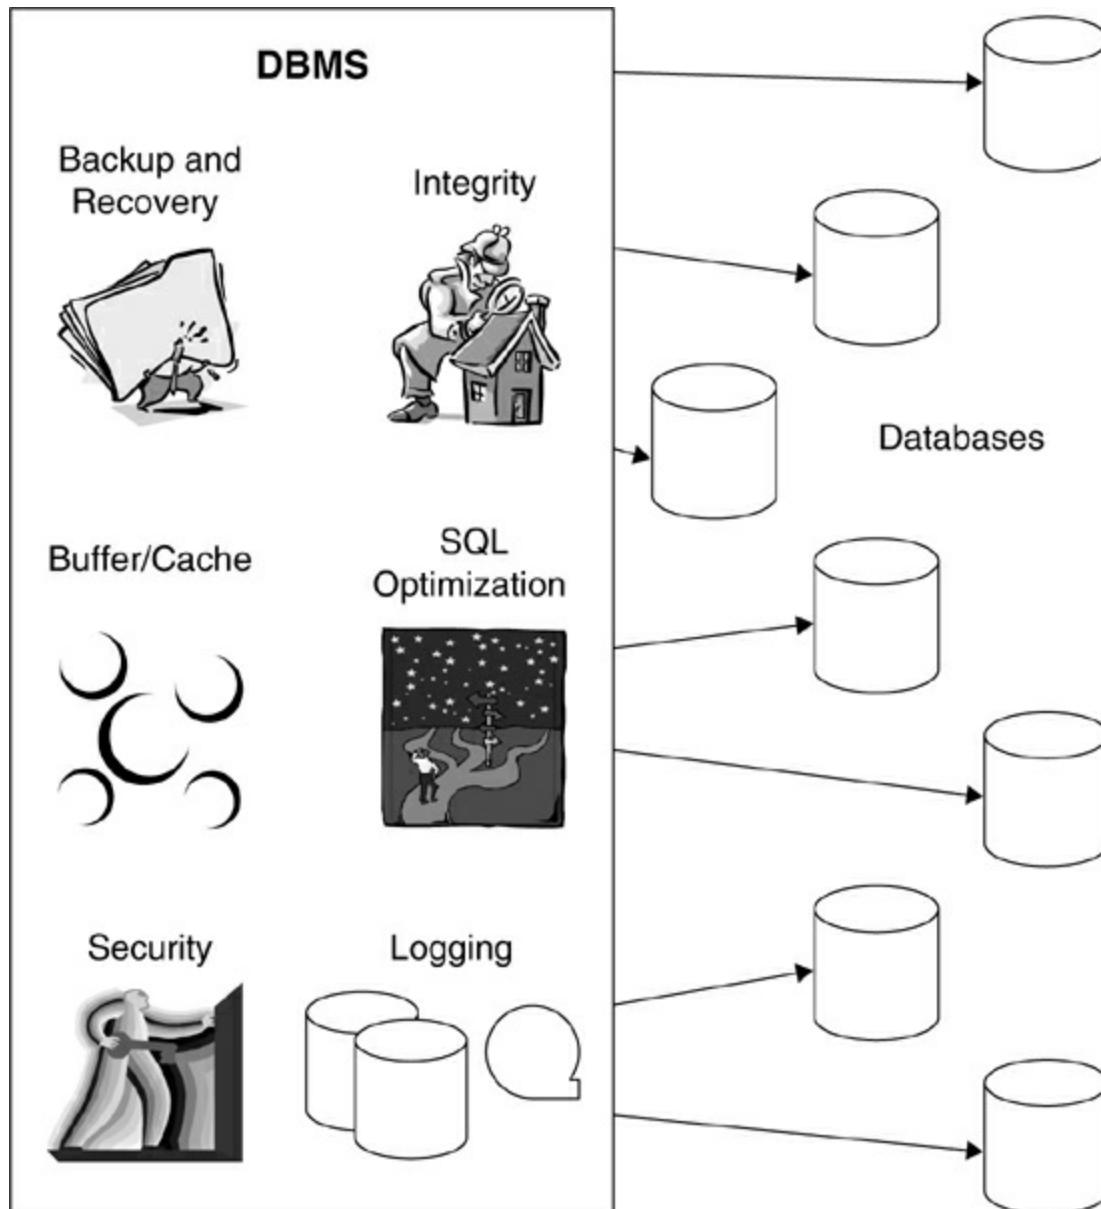


Figure 1.1. DBMS versus database

A database is an organized store of data wherein the data is accessible by named data elements

A DBMS is software that enables end users or application programmers to share data. It provides a systematic method of creating, updating, retrieving,

and storing information in a database. DBMSs also are generally responsible for data integrity, data security, data access control and optimization, automated rollback, restart, and recovery.

In layperson's terms, you can think of a database as a file folder. You can think of the filing cabinet holding the files along with the file labels as the DBMS. A DBMS manages databases. You implement and access database instances using the capabilities of the DBMS. So, DB2 and Oracle and SQL Server are database management systems. Your payroll application uses the payroll database, which may be implemented using DB2 or Oracle or SQL Server.

Why is that important? If we do not use precise terms in the workplace, confusion can result. And confusion leads to over-budget projects, improperly developed systems, and lost productivity.

In addition to database management fundamentals, DBAs must be experts in the specific DBMS products being used, and there may be many in the organization. For example, a large organization may use DB2 on the mainframe, Oracle and Informix on several different UNIX platforms, MySQL on Linux, and SQL Server on Windows. Older legacy systems may use IMS databases, and then there is that one crazy application out there that uses a fringe DBMS like Adabas or Ingres.⁹ And there is also new database technology, such as NoSQL, and column store DBMS offerings like Hadoop, as well as cloud database systems, such as Microsoft's SQL Azure and Google's BigTable.

The DBA group, therefore, must have expertise in each of these different DBMSs and platforms. Furthermore, the DBA must be capable of determining which DBMS and platform are most suited to the needs of each application. This can be a difficult job fraught with politics and conflicting opinions. The DBA group must be able to act as an impartial judge and base implementation decisions on the needs of the application compared to the capabilities and specifications of each DBMS and platform.

DBAs must implement decisions based on the best fit of application, DBMS, and platform.

Once again, for a short introduction to DBMS concepts, refer to [Appendix A](#).

The Management Discipline of Database Administration

Database administration is rarely approached as a management discipline. The term *discipline* implies planning and implementation according to that plan. When database administration is treated as a management discipline, the treatment of data within your organization will improve. It is the difference between being reactive and proactive.

All too frequently the DBA group is overwhelmed by requests and problems. This ensues for many reasons, including understaffing, overcommitment to supporting new (and even old) application development projects, lack of repeatable processes, lack of budget, and on and on. When operating in this manner, the database administrator is being reactive. The reactive DBA functions more like a firefighter. His attention is focused on resolving the biggest problem being brought to his attention. In other words, a reactive DBA attempts to resolve problems only after problems occur.

A proactive DBA can avoid many problems altogether by developing and implementing a strategic blueprint to follow when deploying databases within his organization. This plan should address all phases of the application development life cycle (ADLC) (see [Figure 1.2](#)). The proactive DBA implements practices and procedures to avoid problems before they occur.

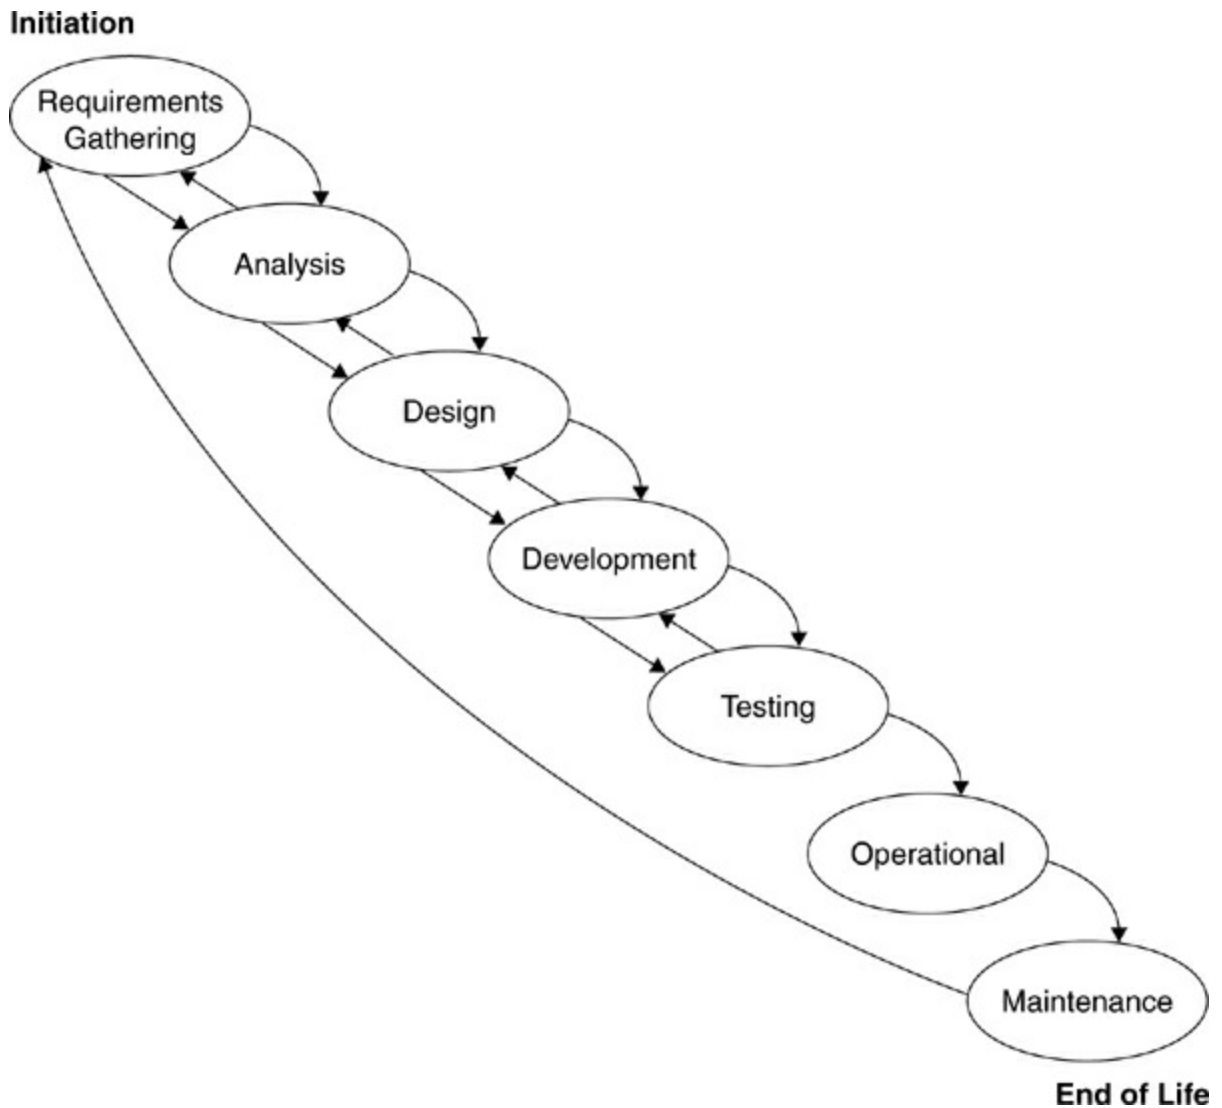


Figure 1.2. The application development life cycle (ADLC)

A proactive DBA can avoid many problems.

A data specialist, usually the DBA, should be involved during each phase of the ADLC. During the initiation and requirements-gathering phase of the project the DBA must be available to identify the data components of the project. He can help to determine if the required data already exists elsewhere in the organization or if the data is brand-new. During the analysis and design phases the rudimentary data requirements must be transformed into a conceptual and logical data model.

Before development can begin, the logical data model must be translated to a physical database design that can be implemented using a DBMS, such as

Oracle or DB2. Sample data must be populated into the physical database to assist with application testing. Furthermore, the DBA must develop and implement a process to refresh test data to enable repeatable test runs.

When the application moves from development to operational status, the DBA must ensure that the DBMS is prepared for the new workload. This preparation includes implementing appropriate security measures, measuring and modifying the storage and memory requirements for the new application, and anticipating the impact of the new workload on existing databases and applications. The DBA is also responsible for migrating the new database from the test environment to production.

While the application is operational, the DBA performs a host of duties, including performance and availability monitoring, tuning, backup and recovery, and managing authorization. But no application or database remains static for long. Because the needs of the business will change, the IT systems that support the business also will change. When maintenance is requested, the DBA becomes engaged in the entire process once again from requirements gathering all the way through to the changes becoming operational. When monitoring reveals a performance shortcoming, the DBA may suggest better-performing alternatives to the development team (if the problem is in the application code) or to the storage team (if the bottleneck is I/O-centric) or recommend other fixes such as new or improved indexes, alternative SQL code, and so forth.

Privacy Policies and Data

Data privacy is a burgeoning issue that is becoming more and more of a burden on organizations, as not only the amount of data under management increases, but the velocity of new data arriving increases.

In this networked world, in which we are thoroughly digitized, with our identities, locations, actions, purchases, associations, movements, and histories stored as so many bits and bytes, we have to ask: Who is collecting all of this? What are they doing with it? With whom are they sharing it? Most of all, individuals are asking, “How can I protect my information from being misused?” These are reasonable questions to ask; we should all want to know the answers.

There are more than 30 federal statutes and over 100 state statutes governing information privacy in the United States. The approach to

protecting privacy has been piecemeal. The European Union, on the other hand, has adopted a Data Protection Directive requiring its member countries to adopt laws that implement its terms. The directive creates rights for persons about whom information is collected, known as “data subjects.” Entities that collect information must give data subjects notice explaining who is collecting the data, who will ultimately have access to it, and why the data is being collected. Data subjects also have the right to access and correct data about them.

There is even a Data Privacy Day,¹⁰ the purpose of which is to celebrate the dignity of the individual expressed through personal information.

The recent bankruptcy of Borders Books provides an example of the impact of privacy policies on corporate data. As part of Borders ceasing operations, Barnes & Noble, a competing bookseller, acquired some of its assets, including Borders’ brand trademarks and its customer list. Their first course of action, however, was to alert the customers as to their rights. The e-mail I received stated (in part):

It’s important for you to understand however you have the absolute right to opt-out of having your customer data transferred to Barnes & Noble. If you would like to opt-out, we will ensure all your data we receive from Borders is disposed of in a secure and confidential manner.

This is one example of how privacy policies can impact the job of database administration and corporate data experts. Of course, you may never work for a company that goes bankrupt, but your company may decide to retire applications and data due to legal regulations, business conditions, or mergers.

Finally, when the application reaches the end of its useful life, the DBA must help to determine the final status of the data used by the application. Is the data no longer required, or do other applications and processes use the data, too? Are there regulations that require the data to be stored longer than the application?¹¹ Does the business have any stated privacy policies that impose special rules for handling the data?

The DBA is responsible for managing the overall database environment. Often this includes installing the DBMS and setting up the IT infrastructure

to allow applications to access databases. These tasks need to be completed before any application programs can be developed. Furthermore, ad hoc database access is a requirement for many organizations.

Additionally, the DBA is in charge of setting up the ad hoc query environment, which includes query and reporting tool evaluation and implementation, establishing policies and procedures to ensure efficient ad hoc queries, and monitoring and tuning ad hoc SQL.

As you can see, a good DBA is integral to the entire application development life cycle. The DBA is “in demand” for his knowledge of data and the way in which data is managed by modern applications.

A good DBA is integral to the entire application development life cycle.

A Day in the Life of a DBA

A day in the life of a DBA is usually quite busy. The DBA is required to maintain production and test environments while at the same time keeping an eye on active application development projects, attending strategy and design meetings, helping to select and evaluate new products, and connecting legacy systems to the Web. And Joe in Accounting, he just submitted that “query from hell” again that is bringing the system to a halt; can you do something about that? All of these things can occur within a single DBA workday.

To add to the chaos, DBAs are expected to know everything about everything. From technical and business jargon to the latest management and technology fads, the DBA is expected to be “in the know.” And don’t expect any private time. A DBA must always be prepared to be interrupted at any time to answer any type of question—and not just about databases, either.

When application problems occur, the database environment is frequently the first thing blamed. The database is “guilty until proven innocent” instead of the other way around. A DBA will never experience an application developer coming to him for help with a question like “I’ve got some really bad SQL here; can you help me fix it?” No, instead the developer will barge into the DBA’s cubicle stating very loudly that “there’s a problem with DB2 [or insert your favorite DBMS here]; why don’t you fix it so my wonderful program can run?”

When application problems occur, the database is “guilty until

proven innocent.”

Thus, the DBA is forced to prove that the database is not the source of the problem. He must know enough about all aspects of information technology to track down errors and exonerate that which is in his realm: the DBMS and database structures he has designed. So he must be an expert in database technology but also know about IT components with which the DBMS interacts: application programming languages, operating systems, network protocols and products, transaction processors, every type of computer hardware imaginable, and more. The need to understand such diverse topics makes the DBA a very valuable resource. It also makes the job interesting and challenging.

And the DBA job is a nonstop job. A DBA must be constantly available to deal with problems, because database applications run around the clock. Most DBAs carry a pager or mobile phone at all times so they can be reached at a moment's notice. If there is a database problem at 2:00 A.M., the DBA must get out of bed, clear his head, and solve the problem to get the applications back up and running. Failure to do so can result in database downtime, and that can completely shut down business processes.

DBAs frequently spend weekends in front of the computer performing database maintenance and reorganizations during off-peak hours. You can't bring mission-critical databases down during the nine-to-five day to maintain them. And frankly, more and more businesses have 24-hour days instead of the mythical eight-hour business days of yore.

If this still sounds intriguing to you, read on. And actually, it isn't as bad as it sounds. The work is interesting, there is always something new to learn, and, as previously mentioned, the pay can be good. The only question is, Can anyone do this type of job for 20 or more years without needing a rest? And, oh, by the way, I think I hear your mobile phone buzzing, so you might want to pause here to see what is wrong.

Evaluating a DBA Job Offer

As a DBA, it is almost inevitable that you will change jobs several times during your career. When making a job change, you will obviously consider requirements such as salary, bonus, benefits, frequency of reviews, and amount of vacation time. However, you also should consider how the

company treats its DBAs. Different organizations place different value on the DBA job. It is imperative for your career development that you scout for progressive organizations that understand the complexity and ongoing learning requirements of the position.

Here are some useful questions to ask:

- Does the company offer regular training for its DBAs to learn new DBMS features and functionality? What about training for related technologies such as programming, networking, e-business, transaction management, message queuing, and the like?
- Does the company allow DBAs to regularly attend local user groups?¹² What about annual user groups at remote locations?
- Are there backup DBAs, or will you be the only one on call 24/7?
- Are there data administration and system administration organizations, or are the DBAs expected to perform all of these duties, too?
- Does the DBA group views its relationship with application development groups as a partnership? Or is the relationship more antagonistic?
- Are DBAs included in design reviews, budgeting discussions, and other high-level IT committees and functions?

The more “yes” answers you get to these questions, the more progressive the DBA environment is.

Database, Data, and System Administration

Some organizations define separate roles for the business aspects of data and the technical aspects of data. The business aspects of data are aligned with a discipline known as data administration, whereas the more technical aspects are handled by database administration. Not every organization has a data administration function. Indeed, many organizations combine data administration into the database administration role.

Many organizations combine data administration into the database administration role.

Sometimes organizations also split up the technical aspects of data management, with the DBA being responsible for using the DBMS and

another role, known as system administration or systems programming, being responsible for installing and upgrading the DBMS.

Data Administration

Data administration (DA) separates the business aspects of data resource management from the technology used to manage data. When the DA function exists in an organization, it is more closely aligned with the actual business users of data. The DA group is responsible for understanding the business lexicon and translating it into a logical data model. Referring back to the application development life cycle, the data administrator (DA) would be involved more in the requirements-gathering, analysis, and design phases; the DBA, in the design, development, testing, and operational phases.

Another differentiating factor between DA and DBA is the focus of their efforts. The DA is responsible for issues such as

- Identifying and cataloging the data required by business users
- Production of conceptual and logical data models to accurately depict the relationship among data elements for business processes
- Production of an enterprise data model that incorporates all of the data used by all of the organization's business processes
- Setting data policies for the organization
- Identifying data owners and stewards
- Setting standards for control and usage of data

In short, the DA can be thought of as the Chief Data Officer of the corporation. The DA position has nothing specifically to do with technology, though. However, in my experience, the DA is never given an executive position. That is actually too bad. Many IT organizations state that they treat data as a corporate asset, but the falsehood of that statement is revealed when you review their actions. Responsibility for data policy is often relegated to technicians who fail to concentrate on the nontechnical, business aspects of data management. Technicians do a good job of ensuring availability, performance, and recoverability but are not usually capable of ensuring data quality and setting corporate policies.

The data administrator can be thought of as the Chief Data Officer of the corporation.

In fact, data is rarely treated as a true corporate asset. Think about the assets that every company has in common: capital, human resources, facilities, and materials. Each of these assets is modeled: chart of accounts, organization charts and reporting hierarchies, building blueprints and office layouts, and bills of materials. Each is tracked and protected. Professional auditors are employed to ensure that no discrepancies exist in our accounting of the assets. Can we say the same thing about data in most organizations?

A mature DA organization is responsible for planning and guiding the data usage requirements throughout the organization. This role encompasses how data is documented, shared, and implemented. A large responsibility of the DA staff is to ensure that data elements are documented properly, usually in a data dictionary or repository. This is another key differentiation between DA and DBA. The DA focuses on the repository, whereas the DBA focuses on the physical databases and DBMS.

Furthermore, the DA deals with metadata, as opposed to the DBA, who deals with data. Metadata is often described as “data about data”; more accurately, metadata is the description of the data and data interfaces required by the business. Data administration is responsible for the business’s metadata strategy.

Examples of metadata include the definition of a data element, business names for a data element, any abbreviations used for that element, and the data type and length of the element. Data without metadata is difficult to use. For example, the number 12 is data, but what kind of data? In other words, what does that 12 mean? Without metadata we have no idea. Consider:

- Is it a date representing December, the twelfth month of the year?
- Or is it a date representing the twelfth day of some month?
- Could it be an age?
- A shoe size?
- Or, heaven forbid, an IQ?
- And so on.

But there are other, more technical aspects of metadata, too. Think about our number 12 again. Consider:

- Is 12 a large number or a small one?
- What is its domain (that is, what is the universe of possible values of

which 12 is but a single value)?

- What about the data type; is it an integer or a decimal number with a 0 scale?

Metadata provides the context in which data can be understood and therefore become information. In many organizations metadata is not methodically captured and cataloged; instead, it exists mostly in the minds of the business users. Where it has been captured in systems it is spread throughout multiple programs in file definitions, documentation in various states of accuracy, or in long-lost program specifications. Some of it, of course, is in the system catalog of the DBMS.

Metadata provides the context in which data can be understood and therefore become information.

A comprehensive metadata strategy will enable an organization to understand the information assets under its control and to measure the value of those assets. Additional coverage of metadata is provided in [Chapter 22](#), “[Metadata Management](#).”

One of the biggest contributions of DA to the corporate data asset is the creation of data models. A conceptual data model outlines data requirements at a very high level. A logical data model provides in-depth details of data types, lengths, relationships, and cardinality. The DA uses normalization techniques to deliver sound data models that accurately depict the data requirements of the organization.

Many DBAs dismiss data administration as mere data modeling, required only because someone needs to talk to those end users to get the database requirements. But a true DA function is much more than mere data modeling. It is a business-oriented management discipline responsible for the data asset of the organization.

Why spend so much time talking about data administration in a book about database administration? Well, few organizations have implemented and staffed a DA role. The larger the organization, the more likely it is that a DA function exists. However, when the DA role is undefined in the organization, the DBA must assume the mantle of data planner and modeler. The DBA usually will not be able to assume all of the functions and responsibility of DA as summarized in this section for the following reasons.

- The DBA has many other technical duties to perform that will consume most of his time.
- The manager of the DBA group typically does not have an executive position that enables him to dictate policy.
- The DBA generally does not have the skills to communicate effectively with business users and build consensus.
- Frankly, most DBAs are happier dealing with technical issues and technicians than with business issues and non-technicians.

When DA and DBA functions exist within the organization, the two groups must work very closely with one another. It is not necessary that both have the same manager, though that could make it easier to facilitate cooperation between DA and DBA. At any rate, it is imperative that there be some degree of cross-pollination of skills between the two groups. The DA will never understand the physical database like a DBA, and the DBA will never understand the business issues of data like a DA, but each job function would be more effective with some knowledge about the other.

Organizations truly concerned about data quality, integrity, and reuse will invariably implement and staff the DA function.

In short, organizations that are truly concerned about data quality, integrity, and reuse will invariably implement and staff the data administration function.

Database Administration

Database administration is the focus of this entire book, so I will not spend a lot of time defining it in this short section. The rest of the book will accomplish that nicely. This section will quickly outline the functions performed by the DBA group when the DA function exists. As illustrated in [Figure 1.3](#), at a high level, the DBA manages data and the DA manages metadata. But let's dig a little deeper.

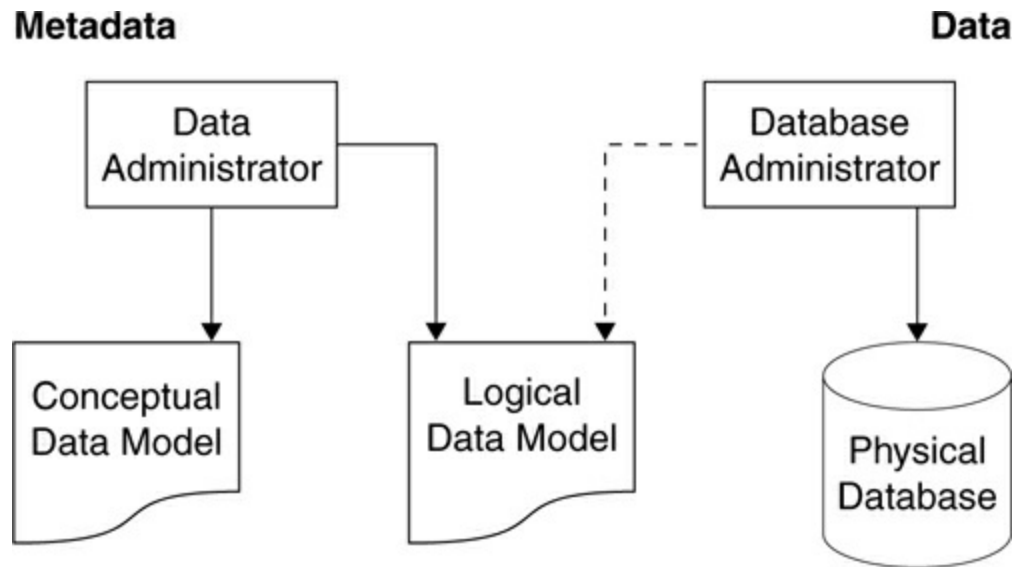


Figure 1.3. DBA versus DA

The first DBA duty is to understand the data models built by DA and to be able to communicate the model to the application developers and other appropriate technicians. The logical data model is the map the DBA will use to create physical databases. The DBA will transform the logical data model into an efficient physical database design. It is essential that the DBA incorporate his knowledge of the DBMS being used to create an efficient and appropriate physical database design from the logical model. The DBA should not rely on the DA for the final physical model any more than the DA should rely on the DBA for the conceptual and logical data models.

The DBA is the conduit for communication between the DA team and the technicians and application programming staff. Of course, the bulk of the DBA’s job is the ongoing support of the databases created from the physical design and the management of the applications that access those databases. An overview of these duties is provided in the upcoming “[DBA Tasks](#)” section of this chapter.

The DBA is the conduit for communication between the DA team and the technicians and application programming staff.

System Administration

Some organizations, once again the larger ones, also have a system administration (SA) or systems programming role that impacts DBMS implementation and operations. When the SA role exists separately from the

DBA role, it is responsible for the installation and setup of the DBMS. The system administrator (SA) typically has no responsibility for database design and support. Instead, the DBA is responsible for the databases and the SA is responsible for DBMS installation, modification, and support. If this distinction is not clear to you, please review [Appendix A](#) for a clear definition of database terminology.

Furthermore, the SA role ensures that the IT infrastructure is implemented such that the DBMS is configured to work with other enabling system software. The SA may need to work with other technicians to configure transaction processors, message queuing software, networking protocols, and operating system parameters to enable the DBMS to operate effectively. The SA ensures that the IT infrastructure is operational for database development by setting up the DBMS appropriately, applying ongoing maintenance from the DBMS vendor, and coordinating migration to new DBMS releases and versions.

The system administrator ensures that the IT infrastructure is operational for database development by setting up the DBMS appropriately, applying ongoing maintenance from the DBMS vendor, and coordinating migration to new DBMS releases and versions.

As with DA, there must be cross-training of skills between the SA and DBA, too. The SA will never understand the physical database like a DBA, but the DBA is unlikely to understand the installation and in-depth technical relationships of system software like the SA. Each job function will be more effective with some knowledge of the other, though.

When an independent SA group does not exist, or when no DBMS-focused SA exists, the DBA assumes responsibility for system administration and programming. The graphic in [Figure 1.4](#) provides a quick delineation of the DA, DBA, and SA tasks.

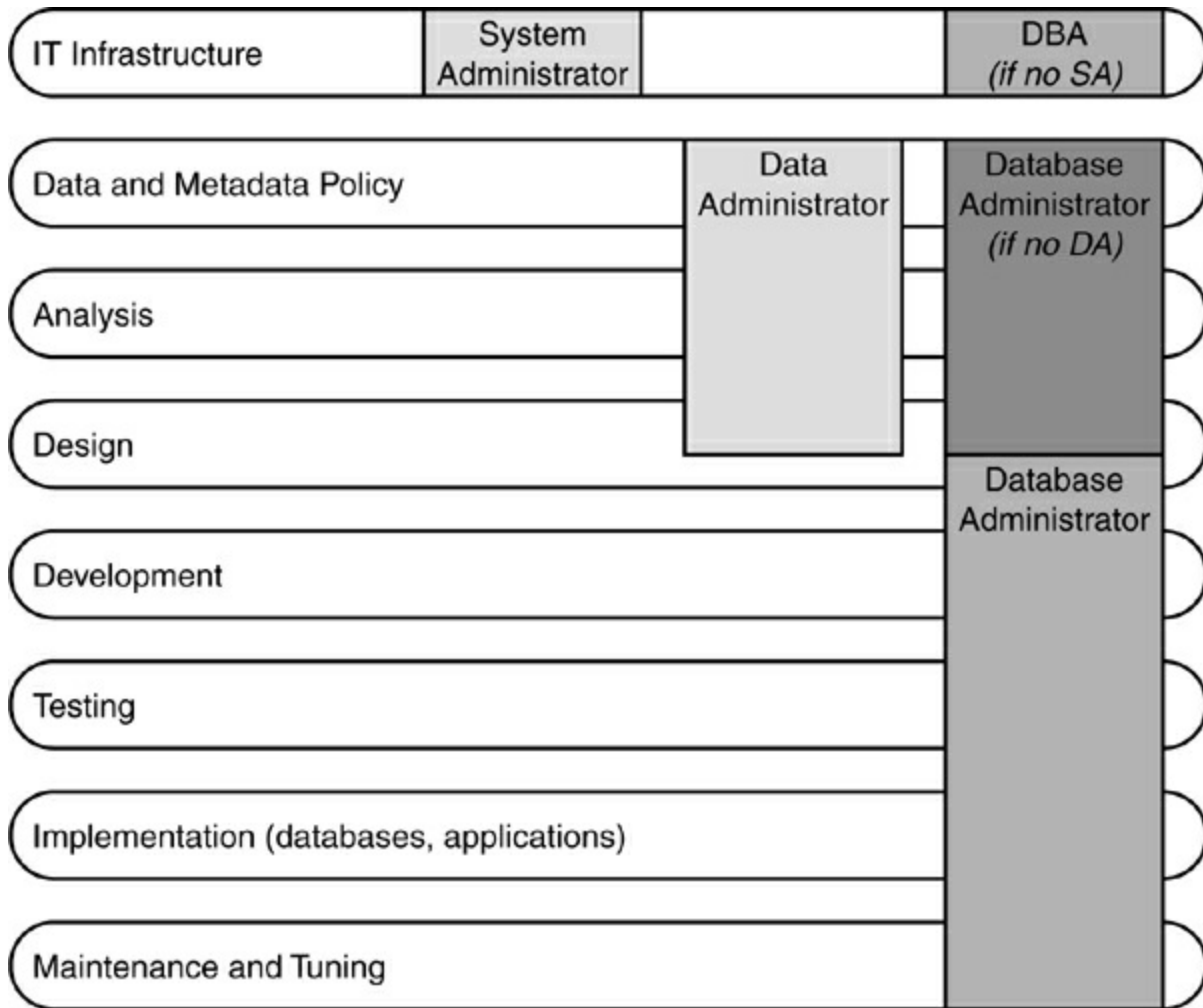


Figure 1.4. DA, DBA, and SA responsibilities

DBA Tasks

A DBA must be capable of performing many tasks to ensure that the organization's data and databases are useful, usable, available, and correct. These tasks include design, performance monitoring and tuning, ensuring availability, authorizing security, backup and recovery, ensuring data integrity, and, really, anything that interfaces with the company's databases. Let's examine each of these topics.

Database Design

The first task that most people think of when they think about DBAs is the ability to create well-designed databases. To properly design and create relational databases DBAs must understand and adhere to sound relational

design practices. They must understand both relational theory and the specific implementation of the RDBMS being used to create the database. Database design requires a sound understanding of conceptual and logical data modeling techniques. The ability to create and interpret entity-relationship diagrams is essential to designing a relational database.

Furthermore, the DBA must be able to transform a logical data model into a physical database implementation. The DBA must ensure that the database design and implementation will enable a useful database for the applications and clients that will use it.

The DBA must be able to transform a logical data model into a physical database.

Indeed, database design is a significant skill for the DBA to possess. However, the job of the DBA often is disproportionately associated with database design. Although designing optimal databases is important, it is a relatively small portion of the DBA's job. A DBA will most likely spend more time administering and tuning databases than in originally designing and building databases.

By no means, though, should you interpret this to mean that database design is not important. A poor relational design can result in poor performance, a database that does not meet the needs of the organization, and potentially inaccurate data.

Performance Monitoring and Tuning

The second role most closely associated with the DBA is performance monitoring and tuning. But what is meant by the term *database performance*? Think, for a moment, of database performance using the familiar concepts of supply and demand. Users demand information from the database. The DBMS supplies information to those requesting it. The rate at which the DBMS supplies the demand for information can be termed *database performance*. But it is not really that simple. Five factors influence database performance: workload, throughput, resources, optimization, and contention.

The [*workload*](#) that is requested of the DBMS defines the demand. It is a combination of online transactions, batch jobs, ad hoc queries, data warehousing and analytical queries, and commands directed through the system at any given time. Workload can fluctuate drastically from day to day,

hour to hour, minute to minute, and even second to second. Sometimes workload can be predicted (such as heavy month-end processing of payroll, or very light access after 7:30 p.m., when most users have left for the day), but at other times it is unpredictable. The overall workload has a major impact on database performance.

Throughput defines the overall capability of the computer hardware and software to process. It is a composite of I/O speed, CPU speed, parallel capabilities of the machine, and the efficiency of the operating system and system software. The hardware and software tools at the disposal of the system are known as the *resources* of the system. Examples include the database kernel, disk space, cache controllers, and microcode.

The fourth defining element of database performance is *optimization*. All types of systems can be optimized, but relational queries are unique in that optimization is primarily accomplished internal to the DBMS. However, there are many other factors that need to be optimized (SQL formulation, database parameters, programming efficiently, and so on) to enable the database optimizer to create the most efficient access paths.

When the demand (workload) for a particular resource is high, [*contention*](#) can result. Contention is the condition in which two or more components of the workload are attempting to use a single resource in a conflicting way (for example, dual updates to the same piece of data). As contention increases, throughput decreases.

Therefore, database performance can be defined as the optimization of resource usage to increase throughput and minimize contention, enabling the largest possible workload to be processed.

Whenever performance problems are encountered by an application that uses a database, the DBA is usually the first one called to resolve the problem. Of course, the DBA cannot manage database performance in a vacuum. Applications regularly communicate with other applications, systems, and components of the IT infrastructure. An effective performance monitoring and tuning strategy requires not just DBMS expertise but knowledge outside the scope of database administration. Many performance management tasks must be shared between the DBA and other technicians. In other words, handling performance problems is truly an enterprise-wide endeavor.

Database performance can be defined as the optimization of resource usage to increase throughput and minimize contention, enabling the largest possible workload to be processed.

The DBA must be vigilant in monitoring system, database, and application performance. As much as possible this should be accomplished using automated software and scripts. Polling system tables and building alerts based on thresholds can be used to proactively identify problems. Alerts can be set up to e-mail the DBA when performance metrics are not within accepted boundaries.

Many tasks and abilities are required of DBAs to ensure efficient access to databases. Some of these abilities include building appropriate indexes, specifying large enough buffers and caches, aligning the database implementation with the IT infrastructure, ongoing monitoring of databases and applications, database reorganization, and adapting to business changes—more users, more data, additional processing, and changing requirements and regulations.

Ensuring Availability

Availability of data and databases is often closely aligned with performance, but it is actually a separate concern. Of course, if the DBMS is offline, performance will be horrible because no data can be accessed. But ensuring database availability is a multifaceted process.

Ensuring database availability is a multi-faceted process.

The first component of availability is keeping the DBMS up and running. Vigilant monitoring and automated alerts can be used to warn of DBMS outages and call for corrective action.

Individual databases also must be maintained such that the data contained therein is available whenever applications and clients require it. Doing so requires the DBA to design the database so that it can be maintained with minimal disruptions, but also to help design applications to minimize conflicts when concurrent access is required.

An additional component of availability is minimizing the amount of downtime required to perform administrative tasks. The faster the DBA can

perform administrative tasks that require databases to be offline, the more available the data becomes. Increasingly, DBMS vendors and independent software vendors (ISVs) are providing nondisruptive utilities that can be performed on databases while applications read and write from them. But these usually require more skill and up-front planning to implement.

The DBA must understand all of these aspects of availability and ensure that each application is receiving the correct level of availability for its needs.

Database Security and Authorization

Once the database is designed and implemented, programmers and users will need to access and modify the data in the database. But only authorized programmers and users should have access to prevent security breaches and improper data modification. It is the responsibility of the DBA to ensure that data is available only to authorized users.

Typically, though not always (see the sidebar “[Centralization of Security](#)”), the DBA works with the internal security features of the DBMS in the form of SQL GRANT and REVOKE statements, as well as any group authorization features of the DBMS. Security must be administered for many actions required by the database environment:

- Creating database objects, including databases, tables, views, and program structures
- Altering the structure of database objects
- Accessing the system catalog
- Reading and modifying data in tables
- Creating and accessing user-defined functions and data types
- Running stored procedures
- Starting and stopping databases and associated database objects
- Setting and modifying DBMS parameters and specifications
- Running database utilities such as LOAD, RECOVER, and REORG

Database security can be enforced in other ways as well. For example, views can be created to block sensitive columns or rows from being viewed by end users and programmers. And the DBA also frequently interfaces with external security methods when they impact database security.

The DBA must understand and be capable of implementing any aspect of

security that impacts access to databases. One area that should be of particular interest given the data breaches in the news these days is SQL injection attacks and how to prevent them.

The DBA must understand the aspects of security that impact access to databases.

Centralization of Security

Some organizations have taken steps to assemble all security and authorization tasks, policies, and procedures in a centralized IT security group. Such an undertaking is not trivial and, as such, it is more common in larger organizations than it is in smaller ones.

Heavily regulated industries may opt to centralize security operations. Furthermore, organizations with mainframe computers tend to consider centralization more often than those without mainframes.

Successfully transferring database security responsibility from the database administration group to a centralized security group requires training the security personnel in database security techniques. Additionally, many of these shops use software that removes security operations from the DBMS and mimics the same functionality in a more traditional security package (such as RACF or ACF2 on the mainframe).

Even taking these issues into consideration, most organizations still rely on the DBA to administer database security.

Governance and Regulatory Compliance

Assuring compliance with industry and governmental regulations is an additional task required of database administration, at least in terms of implementing proper controls. The DBA must work with management, auditors, and business experts to understand the regulations that apply to their industry and the manner in which data is treated.

Certain aspects of regulatory compliance address standard DBA operating procedures. For example, regulations may contain language enforcing specific security and authorization procedures, auditing requirements, data backup specifications, and change management procedures. To ensure compliance, however, stricter documentation may be required, or perhaps a

higher degree of diligence or automation (such as deeper audit tracing).

Other aspects of regulatory compliance may require the DBA to adopt different techniques, tactics, and skills. For example, data retention regulations may require data to be maintained long after it is needed to be stored in a production database, requiring database archiving skills. Or certain data may need to be protected from view, requiring data masking or, in some cases, encryption to be set up.

DBAs should not be tasked with understanding regulations in any depth, nor should they be setting the standards by which the organization complies with the regulations. However, DBAs will get involved in helping to set the proper controls and procedures for compliance projects, specifically and especially with regard to the treatment of data.

DBAs set the proper technological controls and procedures for compliance with regard to the treatment of data.

Backup and Recovery

The DBA must be prepared to recover data in the event of a problem. “Problem” can mean anything from a system glitch or program error to a natural disaster that shuts down an organization. The majority of recoveries today occur as a result of application software error and human error. Hardware failures are not as prevalent as they used to be. In fact, analyst estimates indicate that 80 percent of application errors are due to software failures and human error. The DBA must be prepared to recover data to a usable point, no matter what the cause, and to do so as quickly as possible.

The majority of recoveries today occur as a result of application software error and human error.

The first type of data recovery that usually comes to mind is a *recover to current*, usually in the face of a major shutdown. The end result of the recovery is that the database is brought back to its current state at the time of the failure. Applications are completely unavailable until the recovery is complete.

Another type of traditional recovery is a *point-in-time recovery*. Point-in-time recovery usually is performed to deal with an application-level problem. Conventional techniques to perform a point-in-time recovery will remove the

effects of *all* transactions since a specified point in time. This sometimes can cause problems if there were some valid transactions during that time frame that still need to be applied.

Transaction recovery is a third type of recovery that addresses the shortcomings of the traditional types of recovery: downtime and loss of good data. Thus, transaction recovery is an application recovery whereby the effects of specific transactions during a specified time frame are removed from the database. Therefore, transaction recovery is sometimes referred to as application recovery.

Most technicians think about recovery in order to resolve disasters such as hardware failures. Although hardware failures still occur, and technicians need to be prepared to recover from such failures, most recoveries today are required because of human error or errors in programs.

The DBA must be prepared to deal with all of these types of recovery. This involves developing a backup strategy to ensure that data is not lost in the event of an error in software, hardware, or a manual process. The strategy must be applicable to database processing, so it must include image copies of database files as well as a backup/recovery plan for database logs. It needs to account for any non-database file activity that can impact database applications as well.

The DBA must be prepared to deal with all types of recovery.

Ensuring Data Integrity

A database must be designed to store the correct data in the correct way without that data becoming damaged or corrupted. To ensure this process, the DBA implements integrity rules using features of the DBMS. Three aspects of integrity are relevant to our discussion of databases: physical, semantic, and internal.

Three aspects of integrity are relevant to our discussion of databases: physical, semantic, and internal.

Physical issues can be handled using DBMS features such as domains and data types. The DBA chooses the appropriate data type for each column of each table. This action ensures that only data of that type is stored in the database. That is, the DBMS enforces the integrity of the data with respect to

its type. A column defined as “integer” can contain only integers. Attempts to store nonnumeric or noninteger values in a column defined as integer will fail. DBAs can also use constraints to further delineate the type of data that can be stored in database columns. Most relational DBMS products provide the following types of constraints:

- *Referential constraints* are used to specify the columns that define any relationships between tables. Referential constraints are used to implement referential integrity, which ensures that all intended references from data in one column (or set of columns) of a table are valid with respect to data in another column of the same or a different table.
- *Unique constraints* ensure that the values for a column or a set of columns occur only once in a table.
- *Check constraints* are used to place more complex integrity rules on a column or set of columns in a table. Check constraints are typically defined using SQL and can be used to define the data values that are permissible for a column or set of columns.

Semantic integrity is more difficult to control and less easily defined. DBAs must be prepared to enforce policies and practices to ensure that data stored in their databases is accurate, appropriate, and usable. An example of a semantic issue is the quality of the data in the database. Simply storing any data that meets the physical integrity definitions specified to the database is not enough. Procedures and practices need to be in place to ensure data quality. For example, a customer database that contains a wrong address or phone number for 25 percent of the customers stored therein is an example of a database with poor quality. There is no systematic, physical method of ensuring data accuracy. Data quality is encouraged through proper application code, sound business practices, and specific data policies. Redundancy is another semantic issue. If data elements are stored redundantly throughout the database, the DBA should document this fact and work to ensure that procedures are in place to keep redundant data synchronized and accurate.

The final aspect of integrity is an internal DBMS issue. The DBMS relies upon internal structures and code to maintain links, pointers, and identifiers. In most cases the DBMS will do a good job of maintaining these structures, but the DBA needs to be aware of their existence and how to cope when the

DBMS fails. Internal DBMS integrity is essential in the following areas:

- *Index consistency.* An index is really nothing but an ordered list of pointers to data in database tables. If for some reason the index gets out of sync with the data, indexed access can fail to return the proper data. The DBA has tools at his disposal to check for and remedy these types of errors.
- *Pointer consistency.* Sometimes large multimedia objects are not stored in the same physical files as other data. Therefore, the DBMS requires pointer structures to keep the multimedia data synchronized to the base table data. Once again, these pointers may get out of sync if proper administration procedures are not followed.
- *Backup consistency.* Some DBMS products occasionally take improper backup copies that effectively cannot be used for recovery. It is essential to identify these scenarios and take corrective actions.

Overall, ensuring integrity is an essential DBA skill.

Ensuring integrity is an essential DBA skill.

DBMS Release Migration

The DBA is also responsible for managing the migration from release to release of the DBMS. DBMS products change quite frequently—new versions are usually released every year or so. The task of keeping the DBMS running and up-to-date is an ongoing effort that will consume many DBA cycles. Whatever approach is taken must conform to the needs of the organization, while reducing outages and minimizing the need to change applications.

The task of keeping the DBMS running and up-to-date is an ongoing effort that will consume many DBA cycles.

Jack-of-All-Trades

Databases are at the center of modern applications. If the DBMS fails, applications fail, and if applications fail, the entire business can come to a halt. If databases and applications fail often enough, the entire business can fail. Database administration therefore is critical to the ongoing success of modern business.

Furthermore, databases interact with almost every component of the IT infrastructure. The IT infrastructure of today consists of things such as

- Programming languages and environments such as COBOL, Microsoft Visual Studio, C/C++/C#, Java, and PHP
- Software frameworks such as .NET and J2EE
- Database and process design tools such as ERwin and Rational Rose
- Transaction processing systems such as CICS and Tuxedo
- Application servers such as WebSphere, JBoss, Oracle Application Server, and EA Server
- Message queuing software such as MQSeries and MSMQ
- Networking software and protocols such as SNA, VTAM, and TCP/IP
- Networking hardware such as bridges, routers, hubs, and cabling
- Multiple operating systems such as Windows, z/OS and MVS, UNIX and Linux, and perhaps others
- Data storage hardware and software such as enterprise storage servers, Microsoft SMS, IBM DFHSM, SANs, and NAS
- Operating system security packages such as RACF, ACF2, and Kerberos
- Other types of storage hardware, for example, tape machines, silos, and solid-state (memory-based) storage
- Non-DBMS data set and file storage techniques such as VSAM and b-tree
- NoSQL products such as Hadoop and MongoDB
- Database administration tools and how they interface with other systems management solutions
- Systems management tools and frameworks such as HP OpenView and CA Unicenter
- Operational control software such as batch scheduling software and job entry subsystems
- Software distribution solutions for implementing new versions of system software across the network
- The Internet and Web-enabled databases and applications

- Client/server development techniques (multitier, fat server/thin client, thin server/fat client, etc.)
- Object-oriented and component-based development technologies and techniques such as CORBA, COM, OLE/DB, ADO, and EJB
- Pervasive computing technology devices such as tablets and smartphones

Although it is impossible to become expert in all of these technologies, the DBA should have some knowledge of each of these areas and how they interrelate. Even more importantly, the DBA should have the phone numbers of experts to contact in case any of the associated software and hardware causes database access or performance problems.

The Types of DBAs

There are DBAs who focus on logical design and DBAs who focus on physical design; DBAs who specialize in building systems and DBAs who specialize in maintaining and tuning systems; specialty DBAs and general-purpose DBAs. Truly, the job of DBA encompasses many roles.

Some organizations choose to split DBA responsibilities into separate jobs. Of course, this occurs most frequently in larger organizations, because smaller organizations often cannot afford the luxury of having multiple, specialty DBAs.

Still other companies simply hire DBAs to perform all of the tasks required to design, create, document, tune, and maintain the organization's data, databases, and database management systems. Let's look at some of the more common types of DBA.

System DBA

A *system DBA* focuses on technical rather than business issues, primarily in the system administration area. Typical tasks center on the physical installation and performance of the DBMS software, including

A *system DBA* focuses on technical rather than business issues.

- Installing new DBMS versions and applying maintenance fixes supplied by the DBMS vendor

- Setting and tuning system parameters
- Tuning the operating system, network, and transaction processors to work with the DBMS
- Ensuring appropriate storage for the DBMS
- Enabling the DBMS to work with storage devices and storage management software
- Interfacing with any other technologies required by database applications
- Installing DBA tools and utilities

System DBAs rarely get involved with actual implementation of databases and applications. They may get involved in application tuning efforts when operating system parameters or complex DBMS parameters need to be altered.

Indeed, the job of system DBA usually exists only if the organization does not have an official system administration or systems programming department.

Database Architect

Some organizations designate a separate job, referred to as a *database architect*, for designing and implementing new databases. Database architects are involved in new design and development work only; they do not get involved in maintenance, administration, and tuning efforts for established databases and applications. The database architect designs new databases for new applications or perhaps a new database for an existing application.

The database architect is involved in new design and development work only.

The rationale for creating a separate position is that the skills required for designing new databases are different from the skills required to keep an existing database implementation up and running. A database architect is more likely to have data administration and modeling expertise than a general-purpose DBA, because DA skills are more useful for developing an initial database design.

Typical tasks performed by the database architect include

- Creation of a logical data model (if no DA or data modeler position

exists)

- Translation of logical data models into physical database designs
- Implementing efficient databases, including physical characteristics, index design, and mapping database objects to physical storage devices
- Analysis of data access and modification requirements to ensure efficient SQL and to ensure that the database design is optimal
- Creation of backup and recovery strategies for new databases

Most organizations do not staff a separate database architect position, instead requiring DBAs to work on both new and established database projects.

Database Analyst

Another common staff position is the *database analyst*. There is really no common definition for database analyst. Sometimes junior DBAs are referred to as database analysts. Sometimes a database analyst performs a role similar to the database architect. Sometimes the DA is referred to as the database analyst, or perhaps the data analyst. And sometimes *database analyst* is just another term used by some organizations instead of [*database administrator*](#).

Data Modeler

When the DA role is not defined or staffed, there may be a data modeler role defined. A *data modeler* is usually responsible for a subset of the DA's responsibilities. Data modeling tasks include

- The collection of data requirements for development projects
- Analysis of the data requirements
- Design of project-based conceptual and logical data models
- Creation of a corporate data model and keeping the corporate data model up-to-date
- Working with the DBAs to ensure they have a sound understanding of the data models

Application DBA

In direct contrast to the system DBA is the *application DBA*. Application DBAs focus on database design and the ongoing support and administration of databases for a specific application or applications. The application DBA

is likely to be an expert at writing and debugging complex SQL and understands the best ways to incorporate database requests into application programs. The application DBA also must be capable of performing database change management, performance tuning, and most of the other roles of the DBA. The difference is the focus of the application DBA—not on the overall DBMS implementation and database environment, but on a specific subset of applications (see [Figure 1.5](#)).

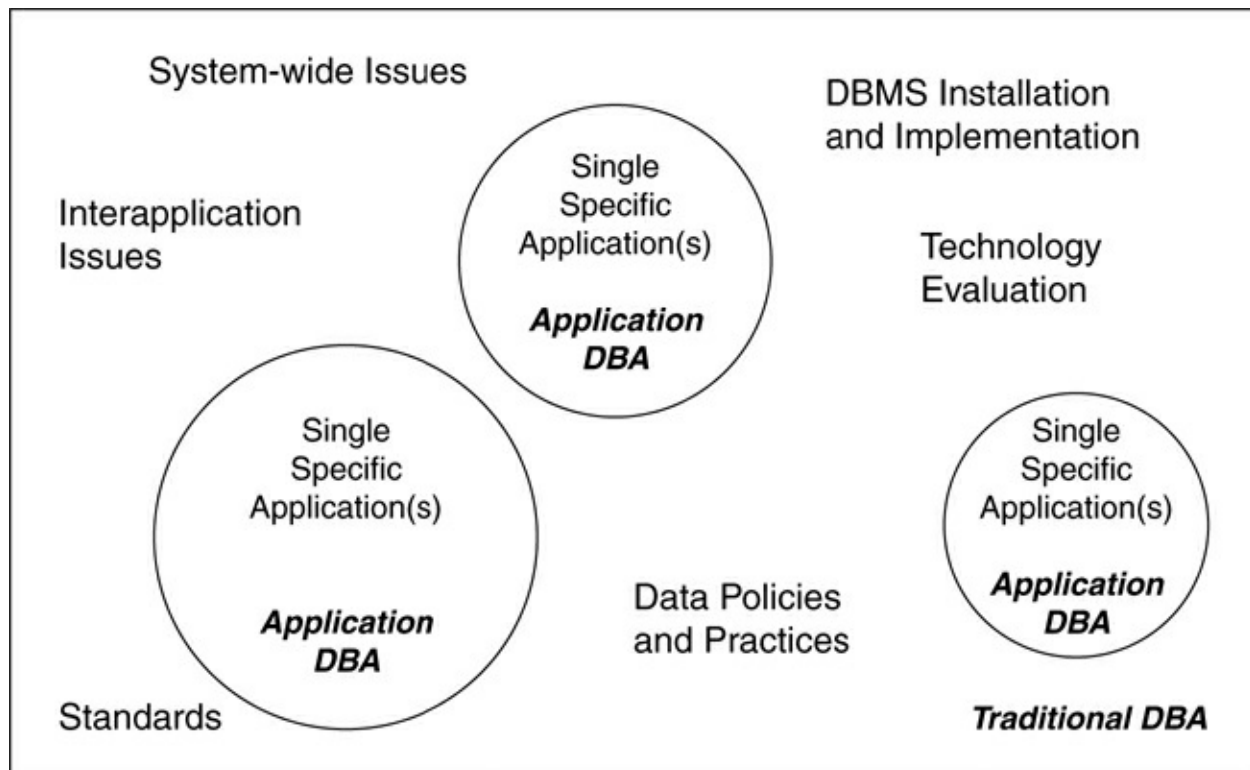


Figure 1.5. Focus of the application DBA

The application DBA focuses on database design and the ongoing support and administration of databases for a specific application or applications.

Not every organization staffs application DBAs. However, when application DBAs exist, general-purpose DBAs are still required to support the overall database environment and infrastructure. When application DBAs do not exist within an organization, general-purpose DBAs are likely to be assigned to support specific applications while also maintaining the organization’s database environment.

There are pros and cons to staffing application DBAs. The arguments in

favor of application DBAs include the following:

- Application DBAs can better focus on an individual application, which can result in better service to the developers of that application.
- The application DBA is more often viewed as an integral component of the development team and therefore is better informed about new development plans and changes to plans.
- Because application DBAs consistently work on a specific set of applications, they can acquire a better overall understanding of how each application works, enabling them to better support the needs of the application developers.
- With a more comprehensive understanding of the application, an application DBA will have a better understanding of how the application impacts the overall business. This knowledge will likely result in the execution of DBA tasks to better support the organization.

But all is not favorable for application DBAs. There are cons to implementing an application DBA role, including these:

- Application DBAs can lose sight of the overall data needs of the organization because of their narrow focus on a single application.
- The application DBA can become isolated. Lack of communication with a centralized DBA group (if one exists) can result in diminished sharing of skills.
- When application DBAs implement useful procedures, it takes more effort to share these procedures among the other DBAs.
- Due to the application-centric nature of application DBAs, they can lose sight of new features and functionality being delivered by the DBMS group.

In general, when staffing application DBAs, be sure to also staff a centralized DBA group. The application DBAs should have primary responsibility for specific applications but should also be viewed as part of the centralized DBA group. Doing so will encourage the pros of implementing application DBAs while discouraging the cons.

When staffing application DBAs, be sure to also staff a centralized DBA group.

Task-Oriented DBA

Larger organizations sometimes create very specialized DBAs who focus on a single specific DBA task. But task-oriented DBAs are quite rare outside of very large IT shops. One example of a task-oriented DBA is a backup and recovery DBA who devotes his entire day to ensuring the recoverability of the organization's databases.

Most organizations cannot afford this level of specialization, but when possible, task-oriented DBAs can ensure that very important DBA tasks are tackled by very knowledgeable specialists.

Performance Analyst

Performance analysts are a specific type of task-oriented DBA. The performance analyst, more common than other task-oriented DBAs, focuses solely on the performance of database applications.

A performance analyst must understand the details and nuances of SQL coding for performance, as well as have the ability to design databases for performance. Performance analysts have knowledge of the DBMS being used at a very detailed technical level so that they can make appropriate changes to DBMS and system parameters when required.

But the performance analyst should not be a system DBA. The performance analyst must be able to speak to application developers in their language in order to help them facilitate the appropriate program changes for performance.

The performance analyst is usually the most skilled, senior member of the DBA staff. It is very likely that a senior DBA grows into this role due to his experience and the respect he has gained in past tuning endeavors.

The performance analyst is usually the most skilled, senior member of the DBA staff.

Data Warehouse Administrator

Organizations that implement data warehouses for performing in-depth data analysis often staff DBAs specifically to monitor and support the data warehouse environment. *Data warehouse administrators* must be capable DBAs, but with a thorough understanding of the differences between a database that supports online transaction processing (OLTP) and a data

warehouse. Common data warehouse administration tasks and requirements include

- Experience with business intelligence, data analytics, and query and reporting tools
- Database design for read-only access
- Data warehousing design issues such as star schema
- Data warehousing technologies such as online analytical processing, or OLAP (including ROLAP, MOLAP, and HOLAP)
- Data transformation and conversion skills
- An understanding of data quality issues
- Experience with data formats for loading and unloading of data
- Middleware implementation and administration

Staffing Considerations

Staffing the DBA organization is not a simple matter. There are several nontrivial considerations that must be addressed, including the size of the DBA staff and the reporting structure for the DBAs.

How Many DBAs?

One of the most difficult things to determine is the optimal number of DBAs required to keep an organization's databases online and operating efficiently. Many organizations try to operate with the minimal number of DBAs on staff, the idea being that having fewer staff members lowers cost. But that assumption may not be true. An overworked DBA staff may make mistakes that cause downtime and operational problems far in excess of the cost of the salary of an additional DBA.

But determining the optimal number of DBAs is not a precise science. It depends on many factors, including

- *Number of databases.* The more databases that need to be supported, the more complex the job of database administration becomes. Each database needs to be designed, implemented, monitored for availability and performance, backed up, and administered. There is a limit to the number of databases that an individual DBA can control.
- *Number of users.* As additional users are brought online as clients of

the applications that access databases, it becomes more difficult to ensure optimal database performance. Additionally, as the number of users increases, the potential for increase in the volume of problems and calls increases, further complicating the DBA's job.

- *Number of applications.* A single database can be used by numerous applications. Indeed, one of the primary benefits of the DBMS is to enable data to be shared across an organization. As more applications are brought online, additional pressure is exerted on the database in terms of performance, availability, and resources required, and more DBAs may be required to support the same number of databases.
- *Service-level agreements (SLAs).* The more restrictive the SLA, the more difficult it becomes for the DBA to deliver the service. For example, an SLA requiring subsecond response time for transactions is more difficult to support than an SLA requiring 3-second response time.
- *Availability requirements.* When databases have an allowable period of scheduled downtime, database administration becomes easier because some DBA tasks either require an outage or are easier when an outage can be taken. Considerations such as supporting e-business transactions and the Web drive the need for 24/7 database availability.
- *Impact of downtime.* The greater the financial impact of a database being unavailable, the more difficult DBA becomes because pressure will be applied to assure greater database availability.
- *Performance requirements.* As the requirements for database access become more performance oriented and faster and more frequent access is dictated, DBA becomes more complicated.
- *Type of applications.* Organizations implement all kinds of applications. The types of applications that must be supported have an impact on the need for DBA services. The DBMS and database needs of a mission-critical application differ from those of a non-mission-critical application. Mission-critical applications are more likely to require constant monitoring and more vigilance to ensure availability. Likewise, OLTP application will have different characteristics and administration requirements from OLAP applications. OLTP transactions are likely to be of shorter duration than OLAP queries; OLTP applications perform both read and write operations, whereas

OLAP applications are usually read only. Each has administration challenges that impose different DBA procedures and needs.

- *Volatility.* The frequency of database change requests is an important factor in the need for additional DBAs. A static database environment requiring few changes will not require the same level of DBA effort as a volatile, frequently changing database environment. Unfortunately, the level of volatility for most databases and applications tends to change dramatically over time. It is difficult to ascertain how volatile an overall database environment will be over its lifetime.
- *DBA staff experience.* The skill of the existing DBA staff will impact whether or not additional DBAs are required. A highly skilled DBA staff will be able to accomplish more than a novice team. Skills more than experience dictate DBA staffing level requirements. A highly motivated DBA with two years of experience might easily outperform a ten-year veteran who is burned out and unmotivated.
- *Programming staff experience.* The less skilled application developers are in database and SQL programming, the more involved DBAs will need to be in the development process, performing tasks such as complex SQL composition, analysis, debugging, tuning, and ensuring connectivity. As the experience of the programming staff increases, the complexity of DBA decreases.
- *End user experience.* When end users access databases directly with ad hoc SQL, their skill level has a direct impact on the complexity of DBA.
- *DBA tools.* DBMS vendors and a number of ISVs offer tools that automate DBA tasks and make administering databases easier. The more tools that are available and the degree to which they are integrated can make DBA tasks less complex. Industry analysts have estimated that without DBA tools up to twice the number of DBAs can be required.

Determining how many DBAs are needed is not a precise science.

The previous list of complexity issues notwithstanding, it is very difficult to combine all of these factors into a formula that will dictate the optimum number of DBAs to employ. Though the research is somewhat dated,

industry analysts at the META Group¹³ created a loose formula for calculating DBA level of effort (LOE). The formula is not a rigorous one but arrives at a DBA LOE by applying weights to six factors: system complexity, application immaturity, end user sophistication, software functionality, system availability, and staff sophistication. By measuring each of these items as much as possible to indicate high or low rates, you plug values into the formula and arrive at a number that is translated into an estimate for the number of DBAs required.

Creating a formula that can dictate the optimal number of DBAs to employ is difficult to achieve.

DBA Reporting Structures

To whom should the DBA group report within the organization? Different companies have taken different approaches to the DBA reporting structure, but a few reporting hierarchies are quite common. There is no one correct answer, but some reporting structures work better than others. Let's review some of the possibilities.

One of the best structures is to create a data resource management (DRM) group that consists of all the data and information specialists of the organization—DA, DBA, data analysts, performance analysts, and so on. This group usually reports directly to the CIO but might report through a systems programming unit, the data center, or technical support. [Figure 1.6](#) depicts a typical reporting structure.

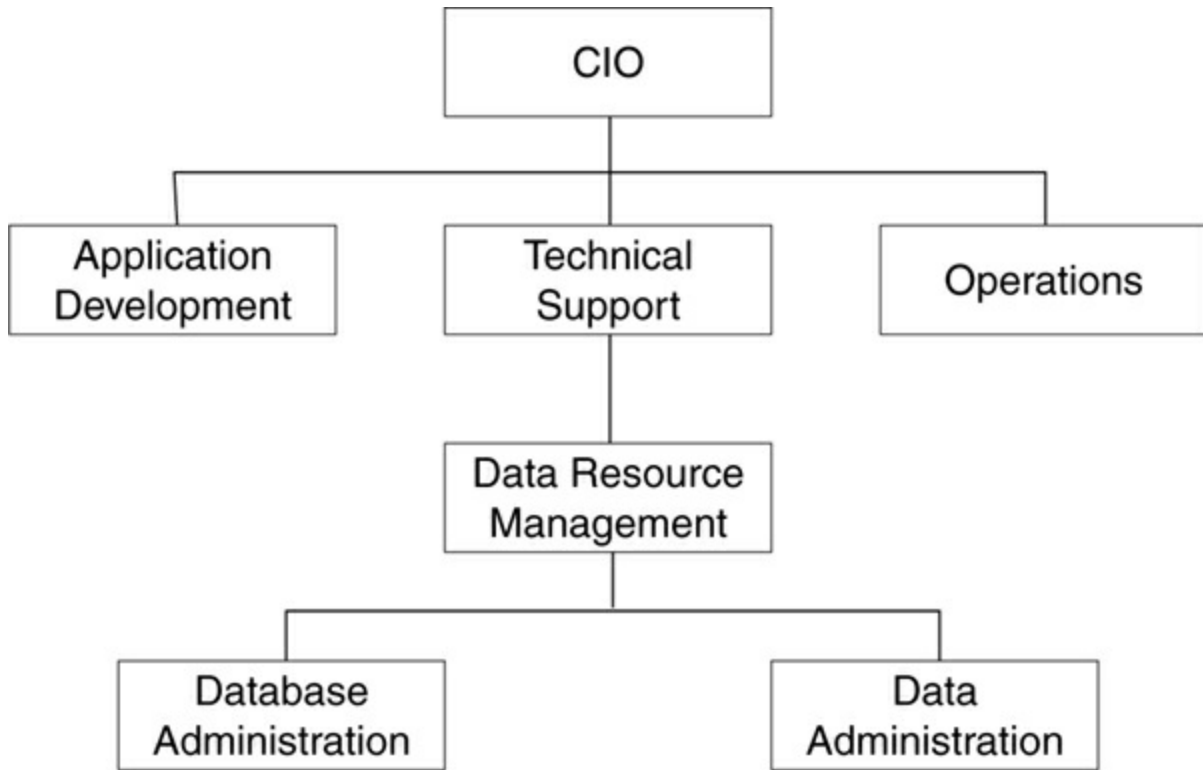


Figure 1.6. Typical DBA reporting structure

When an organization staffs application DBAs, they will be spread out in application groups, typically with a direct reporting line to the business programming managers. Each application development team has a dedicated application DBA resource, as shown in [Figure 1.7](#).

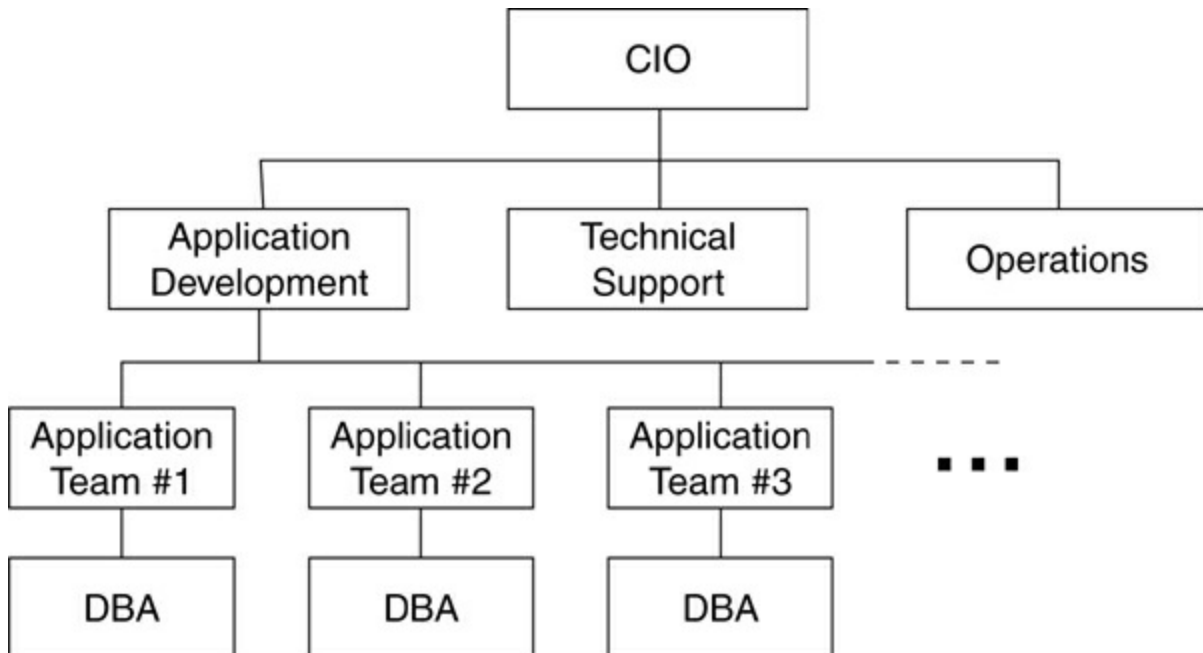


Figure 1.7. Application DBA reporting structure

There are problems with both of these reporting structures, though. The first problem is that DRM should be placed higher in the IT reporting hierarchy. It is a good idea to have the DRM group report directly to the CIO. When an organization understands the importance of data to the health of the organization, placing DRM at this level is encouraged.

Furthermore, when application DBAs exist, they should not report to the application programming manager only. A secondary line of reporting to the DRM group will ensure that DBA skills are shared and communicated throughout the organization. [Figure 1.8](#) delineates the recommended reporting structure for the data resource management group.

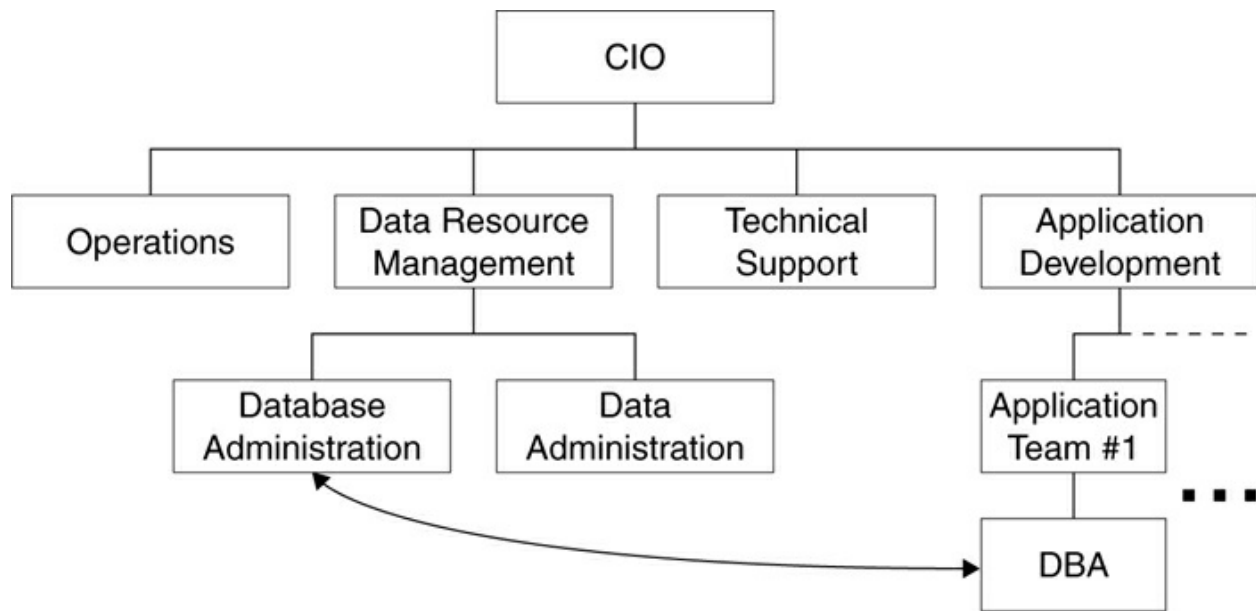


Figure 1.8. Recommended DBA reporting structure

Multiplatform DBA Issues

Managing a multiplatform environment complicates the job of database administration. A whole batch of different problems and issues arise that need to be addressed. The first task is to define the scope of each DBA's job. Does a single DBA administer all of the different DBMSs, or does each DBA focus on supporting only one DBMS?

Managing a multiplatform environment complicates the job of database administration.

This is a particularly thorny issue. On the one hand, the functionality of a DBMS is strikingly similar regardless of platform and vendor. A DBMS is designed to store, retrieve, and protect data. Programmers, programs, and end users all interact with the DBMS to access and modify data. Administration issues are similar—design, creation, optimization, and so on—though each DBMS implements these items differently. So, the case can be made that a DBA should support multiple DBMSs and databases, regardless of platform or vendor.

On the other hand, each DBMS offers different features, functionality, and technology. Keeping all of the differences and nuances straight is a monumental task. Wouldn't it be better to develop platform-expert DBAs? That way, your Oracle DBAs can focus on learning all there is to know about Oracle, your DB2 DBAs can focus on DB2, and so on.

Every organization will have to make this determination based on its particular mix of DBMSs, features, and DBA talent. If your organization uses one DBMS predominantly, with limited use of others, it may make sense for each DBA to support all of them, regardless of platform or vendor. Sparse usage of a DBMS usually means fewer problems and potentially less usage of its more sophisticated features. By tasking your DBAs to be multi-DBMS and multiplatform, you can ensure that the most skilled DBAs in your shop are available for all database administration issues. If your organization uses many different DBMSs, it is probably wise to create specialist DBAs for the heavily used platforms and perhaps share administration duties for the less frequently used platforms among other DBAs.

When DBA duties are shared, be sure to carefully document the skills and knowledge level of each DBA for each DBMS being supported. Take care to set up an effective and fair on-call rotation that does not unduly burden any particular DBA or group of DBAs. Furthermore, use the organizational structure to promote sharing of database standards and procedures across all supported DBMS environments.

Keep in mind, too, that when multiple DBMSs and platforms are supported, you should consider implementing DBA tools, performance monitors, and scripts that can address multiple platforms. For this reason, DBA tools from third-party vendors are usually better for heterogeneous environments than similar tools offered by the DBMS vendors.

When your organization supports multiple DBMSs, the DBA group should

develop guidelines for which DBMS should be used in which situations. These guidelines should not be hard-and-fast rules but instead should provide guidance for the types of applications and databases best supported by each DBMS. Forcing applications into a given DBMS environment is not a good practice. The guidelines should be used simply to assure best fit of application to DBMS. These guidelines should take into account

- Features of each DBMS
- Features and characteristics of the operating system
- Networking capabilities of the DBMS and operating system combination
- DBMS skills of the application developers
- Programming language support
- Any other organizational issues and requirements

Production versus Test

At least two separate environments must be created and supported for a quality database implementation: production and test (or development). New development and maintenance work is performed in the test environment; the operational functioning applications are run in the production environment. It is necessary to completely separate the test environment from the production environment to ensure the integrity and performance of operational work. Failure to separate test and production will cause development activities to impair the day-to-day business of your organization. The last thing you want is for errant program code in the early stages of development to access or modify production data. Test access to production data can cause production performance problems. And, of course, test programs modifying production data can create invalid data.

Separating the test and production environments ensures the integrity and performance of operational work.

The test environment need not be exactly the same as the production environment. The production environment will contain all of the data required to support the operational applications. The test environment, however, can contain only a subset of the data that is required to facilitate acceptable application testing. Furthermore, the test DBMS implementation

usually will not be set up with the same amount of resources as the production environment. For example, less memory will be allocated to buffering and caches, data set allocations will be smaller and on fewer devices, and the DBMS software may be a later version in test than in production (to shake out any bugs in the DBMS code itself before it is trusted to run in production).

The test and production environments should be put together similarly, though. Access to the same system software should be provided in test as is provided to production because the programming staff will need to create their applications in the same type of environment in which they will eventually run.

Another difference in the test environment is the number of databases. Multiple copies of databases may need to be created to support concurrent development by multiple programmers. The DBA must plan and create this environment in such a way as to allow the programming staff to control the contents of the test databases. During the development process programs that modify data in the database may need to be run multiple times. The programmer must be able to ensure, however, that the data at the beginning of each test run is the same. Failure to do so can render the results of the test invalid. Therefore, the DBA must assist the programming staff in the creation of database load and unload jobs to set up test databases for testing runs. Prior to a test run, the database must be loaded with the test data. After the test run, the programmer can examine the output from the program and the contents of the database to determine if the program logic is correct. If not, he can repeat the process, loading to reset the data in the database and retesting. Automated procedures can be put in place to unload the databases impacted by the program and compare the results to the load files.

Trying to predict how test applications will perform once they are moved to production is a difficult task. But the DBA can assist here as well. A relational DBMS typically provides a method to gather statistical information about the contents of its databases. These statistics are then used by the relational optimizer to determine how SQL will retrieve data. This topic is covered in more depth in [Chapter 12, “Application Performance.”](#) But remember, there will be much less data in test databases than in production. In some cases, though, the DBA can set up scripts to read the production statistics and copy them to the test environment, thereby enabling developers

to more accurately gauge how test applications will perform in production.

Some organizations implement more than two environments, as shown in [Figure 1.9](#). If special care is needed for complex application development projects, additional levels of isolated testing may need to occur. For example, a unit test environment may exist for individual program development, then an integration testing environment to ensure that new programs work together, or that new programs work correctly with existing programs. A quality assurance (QA) environment may need to be established to perform rigorous testing against new and modified programs before they are migrated to the production environment.

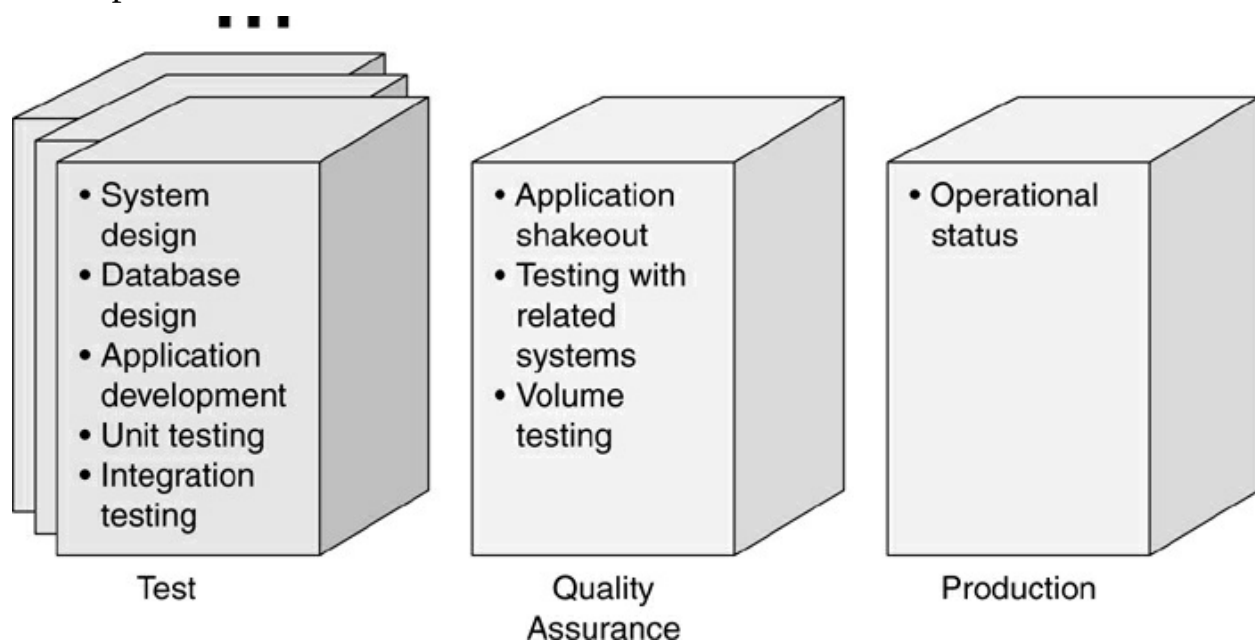


Figure 1.9. Establishing multiple database environments

A QA environment may be needed to perform rigorous testing against new and modified programs before they are migrated.

The Impact of Newer Technology on DBA

The DBA is at the center of the action whenever new ways of doing business and new technologies are introduced to the organization. Data is at the heart of any application, and most new technology impacts data as it is adopted by application developers. Indeed, data is the lifeblood of modern business; the database houses the data; and the DBA is the expert who understands database technology—and in particular, how databases can be integrated with

other new technologies.

Let's examine three specific newer technologies that rely on database administration, at least somewhat, to be effectively implemented: database-coupled application logic, Internet-enabled e-business development, and handheld computing.

Procedural DBAs: Managing Database Logic

Traditionally, the domain of a database management system was, appropriately enough, to store, manage, and access data. Although these core capabilities are still required of modern DBMS products, additional procedural functionality is slowly becoming not just a nice feature to have, but a necessity. Features such as triggers, user-defined functions, and stored procedures provide the ability to define business rules to the DBMS instead of in separate application programs. These features tightly couple application logic to the database server.

Since all of the most popular RDBMS products provide sometimes complex features to facilitate database-coupled procedural logic, additional work is required to manage and ensure the optimal usage of these features. This requires an expansion of the management discipline of database administration. Typically, as new features are added, the administration, design, and management of these features are assigned to the DBA by default. Without proper planning and preparation, this can lead to chaos. But first let's quickly examine how database logic is stored in a DBMS.

Stored Procedures

Stored procedures can be thought of as programs that live in a database. The procedural logic of a stored procedure is maintained, administered, and executed through the database commands. The primary reason for using stored procedures is to move application code from a client workstation to the database server. Stored procedures typically consume less overhead in a client/server environment because one client can invoke a stored procedure that causes multiple SQL statements to be run. The alternative, the client executing multiple SQL statements directly, is less efficient because it increases network traffic which can degrade overall application performance. A stored procedure is a freestanding database object; it is not "physically" associated with any other object in the database. A stored procedure can access and/or modify data in many tables.

Triggers

Triggers are event-driven specialized procedures that are attached to database tables. The trigger code is executed by the RDBMS automatically as data changes in the database. Each trigger is attached to a single, specified table. Triggers can be thought of as an advanced form of rule or constraint written using procedural logic. A trigger cannot be directly called or executed; it is automatically executed (or “fired”) by the RDBMS as the result of an INSERT, UPDATE, or DELETE SQL statement being issued on its associated table. Once a trigger is created, it is always executed when its “firing” event occurs.

User-Defined Functions

A UDF, or user-defined function, provides a result based upon a set of input values. UDFs are programs that can be executed in place of standard, built-in SQL scalar or column functions. A scalar function transforms data for each row of a result set; a column function evaluates each value for a particular column in each row of the results set and returns a single value. Once written, and defined to the RDBMS, a UDF becomes available just like any other built-in database function.

[Table 1.2](#) summarizes the differences among stored procedures, triggers, and user-defined functions.

Table 1.2. Procedural Database Objects

Object Type	Definition	Executed	How
Stored procedures	Program logic executed on the database server	By request	Explicit
Triggers	Event-driven procedures attached to database tables	Automatic	Implicit
UDFs	Program logic extending SQL functionality	By request in SQL	Explicit

Administering Stored Procedures, Triggers, and UDFs

Once applications and developers begin to rely upon stored procedures, triggers, and UDFs, steps need to be taken to ensure they are managed properly. DBAs must grapple with the issues of quality, maintainability, efficiency, and availability. How and when will these procedural objects be

tested? The impact of a failure is enterprise-wide, not relegated to a single application. This increases the visibility and criticality of these objects. Who is responsible if they fail? The answer must be: a DBA.

The role of administering procedural database logic should fall upon someone skilled in that discipline. A new type of DBA is required to accommodate the administration of database procedural logic. This new role can be defined as a *procedural DBA*.

The procedural DBA should be responsible for those database management activities that require procedural logic support. This should include primary responsibility for ensuring that stored procedures, triggers, and user-defined functions are effectively planned, implemented, shared, and reused. The procedural DBA also should take primary responsibility for coding and testing all triggers. Stored procedures and user-defined functions, however, will most likely be coded by application programmers and reviewed for accuracy and performance by procedural DBAs.

The procedural DBA should participate in and lead the review and administration of all procedural database objects: that is, triggers, stored procedures, and UDFs. Although procedural DBAs are unlikely to be as skilled at programming as application programmers or systems analysts, they must be able to write and review program code reasonably well. The skill level required depends on what languages are supported by the DBMS for creating procedural objects, the rate and level of adoption within the organization, and whether or not an internal organization exists for creating common, reusable programs. [Table 1.3](#) provides a reasonable level of procedural DBA involvement for each type of procedural object. Additionally, the procedural DBA should be on call for any problems that occur to database procedural objects in production.

Table 1.3. Procedural DBA Involvement by Object

Object Type	Level of Procedural DBA Involvement
Stored procedures	Not likely to actually write stored procedures; must review all code before migration to production; should communicate availability and promote reuse
Triggers	Likely to actually write, test, and debug triggers; must communicate deployment of triggers to ensure application awareness
UDFs	Not likely to actually write user-defined functions; should work closely with the development team to ensure UDF functionality and performance; must review all code before migration to production; must communicate availability and promote reuse

The procedural DBA should participate in and lead the review and administration of all procedural database objects: that is, triggers, stored procedures, and UDFs.

The role of the procedural DBA requires communication skills as much as it requires technological acumen (see [Figure 1.10](#)). In addition to managing and optimizing database procedural objects, the procedural DBA must inform the development community of new triggers, stored procedures, and UDFs. Furthermore, the DBA must promote reuse. If the programmers do not know that these objects exist, they will never be used. Other procedural administrative functions can be allocated to the procedural DBA. Depending upon the number of DBAs and the amount of application development being done, the procedural DBA can be assigned to additional functions, such as

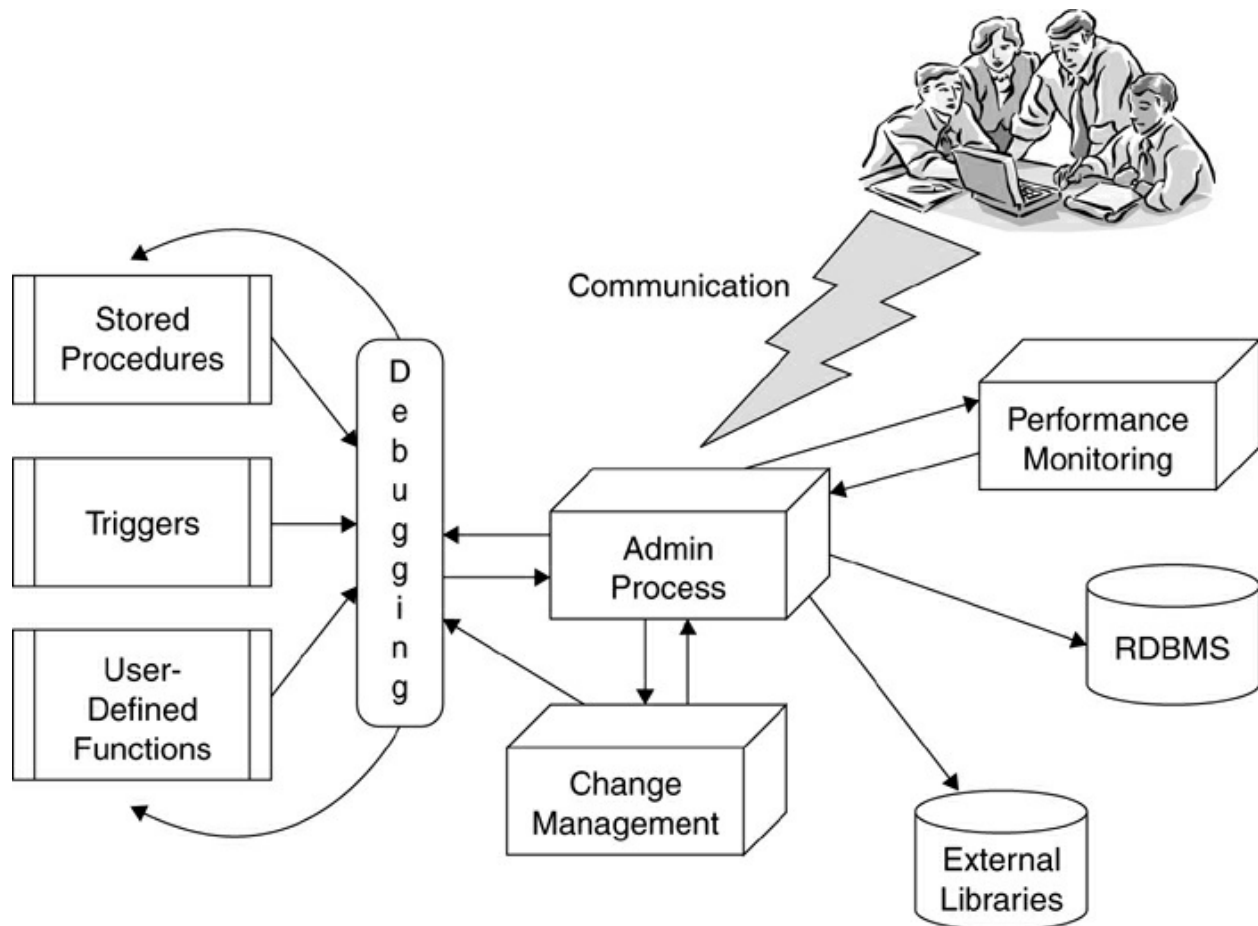


Figure 1.10. Procedural DBA duties

- Participating in application code design reviews
- Reviewing and analyzing SQL access paths (from EXPLAIN or SHOWPLAN)
- Debugging SQL
- Writing and analyzing complex SQL statements
- Rewriting queries for optimal execution

Off-loading coding-related tasks to the procedural DBA can help the other staff DBAs to concentrate on the actual physical design and implementation of databases, resulting in much better-designed databases. The procedural DBA should still report through the same management unit as the traditional DBA. Doing so enables better sharing of skills between the procedural DBAs and traditional data-focused DBAs. Of course, there will need to be a greater synergy between the procedural DBA and the application programmers. The typical job path for the procedural DBA should be from the application

programming ranks because this is where the coding skill base exists.

The Internet: From DBA to eDBA

Although at this late date the Internet can hardly be considered a trend, companies and technologists are still adapting their processes to align with e-commerce. Organizations of every type and size are using Internet technologies to speed up business processes, and database administration practices and procedures are impacted by the adoption of Internet-enabled applications and databases.

E-businesses must be able to adapt and react to constant change. When your business is online, it never closes. People expect full functionality on Web sites they visit regardless of the time. And the Web *is* worldwide. It may be two o'clock in the morning in New York City, but it is always prime time somewhere in the world. An e-business must be available and prepared to engage with customers 24 hours a day, 365 days a year (366 during leap years). Failing to do so can cause loss of business. When a Web site is down, the customer will go elsewhere to do business because the competition is just a simple mouse click away. So those who manage an e-business must be adept, proactive, and ever vigilant.

Adopting the frantic pace of an e-business requires changes for those who keep the e-business operational. DBAs are extremely impacted by e-business. The need to integrate the Web with traditional IT services, such as the DBMS, places high expectations on database administrators.

A DBA who is capable of managing Web-based applications because he understands the special issues that arise because of the Internet is an eDBA. An eDBA also needs to have all of the knowledge and training of a traditional DBA, but these skills need to be adapted to suit applications and databases that are Internet enabled. When the Web is coupled with traditional applications and databases, a complex infrastructure is the result (see [Figure 1.11](#)). The eDBA must be capable of navigating this complex, heterogeneous infrastructure and providing expertise wherever databases interact within this infrastructure.

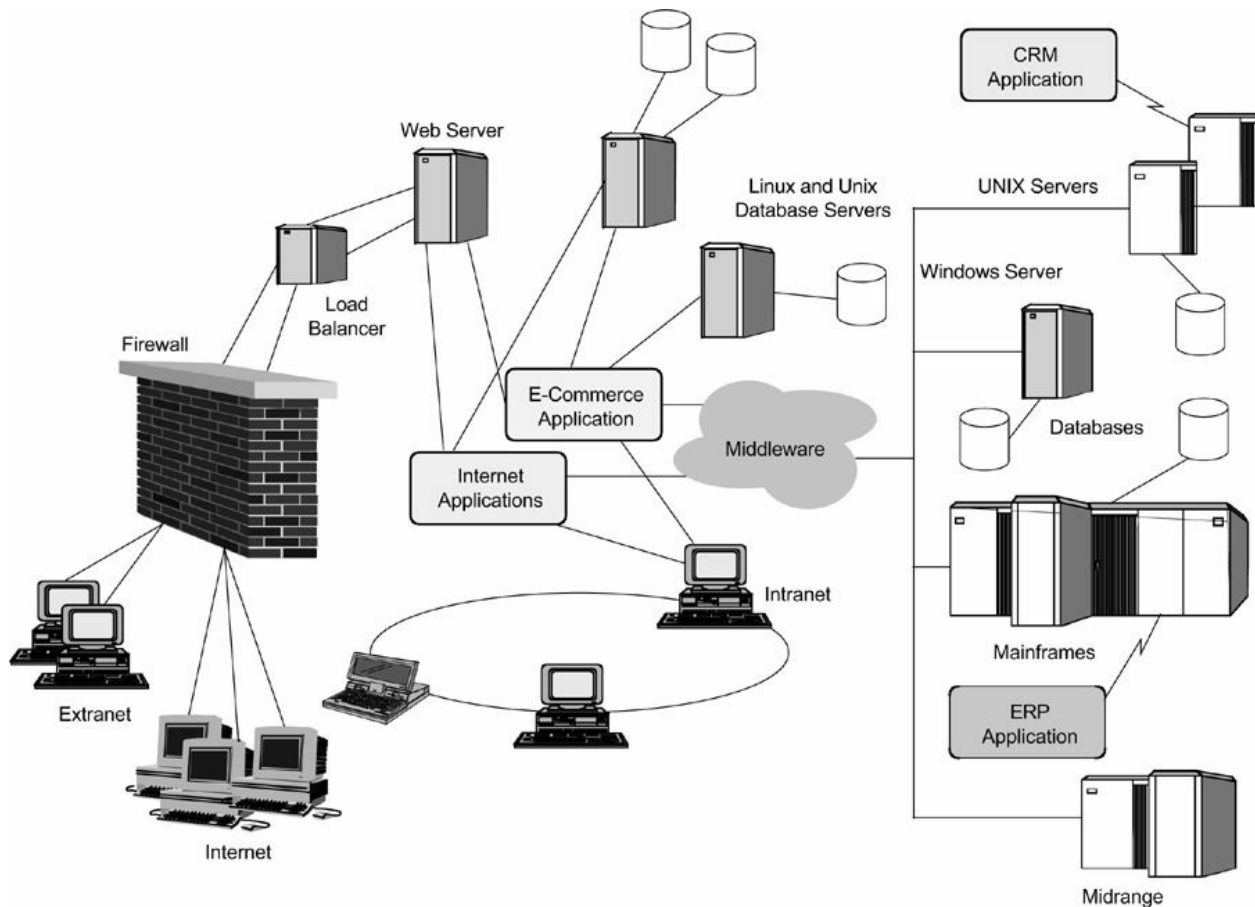


Figure 1.11. The complex infrastructure resulting from Web-to-database capabilities

A DBA who is capable of managing Web-based applications because he understands the special issues that arise because of the Internet is an eDBA.

Indeed, there are many factors that impact database administration when the Internet is coupled with database technology. Some of these issues include

- 24/7 data availability
- Adoption of new technologies such as Java and XML
- Connectivity to the Web
- Integrating legacy data with modern Web-based applications
- Database and application architecture
- Web-based administration

- Performance engineering for the Internet
- Unpredictable workload

The Personal DBA and the Cloud

Personal devices, usually smartphones but also PDAs (personal digital assistants), are fast becoming a necessity for modern executives and businessmen. A smartphone is a handheld computing device. And sometimes it will have a database management system running on it. Why is that interesting? Does it change the way you will use your device? What will that mean to your IT department?

Popular mobile computing platforms include Symbian OS, Windows Mobile, iPhone OS, and Android.

Smartphones offer many benefits. The devices are small and therefore easily transportable. They do not interfere with a mobile worker's ability to be mobile. And because most everyone carries a mobile phone these days, boosting its capabilities to perform computing tasks is a no-brainer.

Perhaps the biggest benefit of these devices is their ability to run mobile applications. It is not uncommon for an enterprise mobile application to rely on information retrieved from a main computer server to be delivered to the mobile device pervasively, at any place and at any time it is required. Mobile applications using the cloud as the back end are becoming more and more popular.

Enterprise mobile applications rely on information retrieved from a main computer server.

What Is Cloud Computing?

Cloud computing offers a new model for the delivery of IT resources to users. The primary defining characteristic of cloud computing is to give the illusion of on-demand access to an infinite amount of computing resources. A good example of a cloud computing service is offered by Salesforce.com, which delivers access to a CRM application over the Web.

Another aspect prevalent with cloud computing offerings is that users can rent computing power with no commitment. Instead of buying a server, you can rent the use of one and pay just for what you use. This used to be referred to as utility computing because it mimics how people

pay for utilities, such as water or electricity. It is a “pay as you go” service.

From a database perspective, there are several cloud offerings, including Amazon’s SimpleDB and Google’s App Engine Datastore. Additionally, Microsoft’s SQL Azure supports data in the cloud.

Cloud computing can allow even the smallest of organizations—or indeed, a single individual—to obtain computing resources in ways not previously possible.

The design and implementation of mobile applications is not as straightforward as desktop PC application development. It is very important for mobile application developers to consider the context in which the application will be used.

Although the benefits are significant, there are challenges to be faced as organizations incorporate personal devices into their infrastructure. The data on the device must be managed professionally to ensure integrity and reliability. Because the device is remote, sharing of data can be difficult. The business data on a smartphone must be reliably synchronized with existing enterprise systems and databases. And from a business and compliance perspective, it can be difficult to assess the risk associated with a mobile application. Mobile devices are easily misplaced and without proper security can result in a data breach.

All of the major DBMS vendors provide small-footprint versions of their flagship products to run on personal devices. For example, IBM markets DB2 Personal Edition, Oracle sells Oracle Database Lite, Microsoft provides SQL Server Compact, and Sybase offers Adaptive Server Anywhere. The general idea is to store a small amount of critical data on the mobile device in a database. The local database is later synchronized to long-term data stores on enterprise database servers. Each mobile DBMS provides technology to synchronize data back and forth from the mobile device to the enterprise server platforms.

Impact on DBAs

DBAs are not needed to manage and work on the database on each PDA, but the job of the DBA will be impacted by this development. A database the size of those stored on PDAs should not require the in-depth tuning and

administration that are required of enterprise database implementations. However, DBAs will be called upon to help design appropriately implemented databases for small-form-factor devices like PDAs. But this is not the biggest impact.

A big impact on the DBA is in the planning for and management of the data synchronization from hundreds or thousands of PDAs. When should synchronization be scheduled? How will it impact applications that use large production databases that are involved in the synchronization? How can you ensure that mobile users will synchronize their data reliably and on schedule?

Additionally, for cloud implementations, DBAs will likely be responsible for assuring reliable data availability. Designing and tuning a database implementation for cloud computing can require significant resources to manage large amounts of data and assure around-the-clock availability.

These are not minor issues. Before implementing a large battalion of mobile database users who must synchronize their data, make sure that your DBA staff is prepared for the impact on their databases. Like most anything, failure to prepare is a recipe for sure disaster. But prepare we must. The DBA staff must be ready to support the mobile workforce by understanding data synchronization technology, cloud computing, and the potential need for remote database users in your organization.

An additional step in the preparation process for supporting databases on handheld devices is to review the applications in your organization and try to determine which might be impacted first. Companies with remote workers such as a distributed sales force or delivery tracking services will most likely be the first impacted. Take some time to review the data requirements of those applications and how a large influx of remote connections might impact the current systems.

Pervasive computing and the mobile workforce are here to stay. And the DBA staff must be ready to support these mobile workers with a valid, shared data infrastructure.

NoSQL, Big Data, and the DBA

NoSQL is another trend that can impact the job performed by the DBA. NoSQL is a movement that, at its most basic, is described by its title. That is, a NoSQL DBMS does not support SQL. At a high level, NoSQL implies nonrelational, distributed, flexible, and scalable. Most NoSQL offerings are

also open source.

NoSQL grew out of the perceived need for “modern” database systems to support Web initiatives. Additionally, some common attributes of NoSQL DBMSs include the lack of a schema, simplicity of use, replication support, and an “eventually consistent” capability (instead of the typical ACID—atomicity, consistency, isolation, and durability—transaction capability). It really does not mean no SQL support. Some NoSQL offerings have begun to support SQL, leading some pundits to define NoSQL as “Not Only SQL.”

Examples of NoSQL Offerings

- Cassandra: <http://cassandra.apache.org/>
 - CouchDB: <http://couchdb.apache.org/>
 - HBase: <http://hbase.apache.org/>
 - mongoDB: www.mongodb.org/
 - Riak: www.basho.com/
-

The NoSQL movement is tied to the Big Data movement. NoSQL databases are designed to deliver low-cost storage and access to large amounts of data.

With NoSQL implementations the data is typically accessed in one way and there is little to no flexibility in terms of generating off-the-cuff, ad hoc queries.

Impact on DBAs

Another hallmark of NoSQL is that it requires very little database administration work. Of course, the DBMS must be set up and managed, and the data must be backed up. So be careful when anyone tries to tell you that any database system requires no database administration. It just isn't true.

New Technology Impacts on DBA

As new technology is introduced to the organization, the DBA group is typically the first to examine and use it. The preceding technologies are merely examples of recent trends and technologies that require database administration for efficient and effective implementation. Most new technologies will have some impact on the role of the DBA.

DBA Certification

Professional certification is an ongoing trend in IT and is available for many different IT jobs. The availability and levels of certification have been growing at an alarming rate for database administration. Certification programs are available for most of the popular DBMS platforms, including IBM DB2, Microsoft SQL Server, and Oracle. The concept behind certification of DBAs is to certify that an individual is capable of performing database administration tasks and duties.

This is a noble goal, but the problem is that passing a test is not a viable indicator of being able to perform a complex job like DBA. Some things you just have to learn by doing. Now I am not saying that certification is useless. Indeed, taking the test and focusing on the questions you miss can help to point out areas of weakness upon which you can improve. But does anyone *really* believe that someone who passes a formal test will be as capable as someone with several years of experience as a DBA? Organizations should hire DBAs based on past experience that indicates a level of capability. Of course, someone with both experience and certification is better than someone with only one of the two.

Certification can make you more employable.

I do recommend that professional DBAs take the time to study and pass the certification exams, not because certification will make you a better DBA, but because it will make you more employable. Some companies will hire only certified professionals. The trend toward using certification to guide hiring practices will increase because of increasing IT complexity. If you think you might change jobs at some point in your career (and who among us will not?), certification is a worthwhile pursuit.

Keep in mind that the DBA certification tests sometimes ask arcane syntax questions that are not really good indicators of a DBA's skills. Getting the syntax 100 percent accurate is what manuals and DBA tools are for. There is no reason to memorize syntax because it tends to change quite often. It is better to know where to find the syntax, parameters, and answers to your questions when you need them, that is, which manuals and textbooks contain the needed information. DBAs should possess a broad overarching knowledge of DBMS concepts and IT fundamentals, and a good knowledge

of the way in which their organization's database systems work. Memorizing every detail about SQL syntax and structure is a waste of time because it is complex and changes all the time. In other words, it is better to know off the top of your head that something can (or cannot) be done than to know the exact syntax for how to accomplish it.

Certification tests sometimes ask arcane syntax questions that are not really good indicators of a DBA's skills.

If you decide to pursue certification, take the time to prepare for the tests. There are books and self-learning software titles available that can be quite useful. These books and programs cover the most likely test topics and provide sample questions to help you prepare. In many ways it is like preparing for a college entrance exam, such as the SATs.

And once you earn your certification, make sure you display it proudly on your résumé and your business card (if your company allows it).

Consult [Table 1.4](#) for Web sites that contain information about professional certification for the most popular DBMS products.

Table 1.4. Sources of DBA Certification Information

DBMS	Web Site with Certification Information
Oracle, MySQL	http://education.oracle.com/
Microsoft SQL Server	www.microsoft.com/learning/en/us/default.aspx
IBM DB2, Informix	www.ibm.com/certify
Sybase	www.sybase.com/detail?id=1009636

The Rest of the Book

This first chapter has introduced you to the world of the DBA. Ideally you have garnered a respect for the complexity of the position and the qualities required of a good DBA. The remainder of the book will examine the details of the tasks, roles, and responsibilities required of the DBA. Read on to learn of the challenges faced by DBAs and how to overcome them.

Review

1. At a high level, discuss the primary job responsibilities of a DBA.

2. What is the single biggest problem faced by organizations using relational databases?
3. What is the difference between a data administrator and a database administrator?
4. What factors determine the number of DBAs required to properly support an organization's database environment?
5. How does new technology impact the job of the DBA?
6. Cite the technology influences that mandate the need for procedural DBAs.
7. What is the difference between a database architect and a system administrator?
8. What job function is most likely to be responsible for installing a new DBMS release?
9. What are the three types of integrity that DBAs must understand?
10. Is a certified DBA necessarily a qualified DBA? Why or why not?

Bonus Question

Why must the DBA be prepared to function as a jack-of-all-trades?

2. Creating the Database Environment

One of the primary tasks associated with the job of DBA is the process of choosing and installing a DBMS. Unfortunately, many business executives and IT professionals without database management background assume that once the DBMS is installed, the bulk of the work is done. The truth is, choosing and installing the DBMS is hardly the most difficult part of a DBA's job. Establishing a usable database environment requires a great deal of skill, knowledge, and consideration. This chapter will outline the principles involved in establishing a usable database environment.

Defining the Organization's DBMS Strategy

The process of choosing a suitable DBMS for enterprise database management is not as difficult as it used to be. The number of major DBMS vendors has dwindled due to industry consolidation and domination of the sector by a few very large players.

Choosing a suitable DBMS for enterprise database management is not as difficult as it used to be.

Yet, large and medium-size organizations typically run multiple DBMS products, from as few as two to as many as ten. For example, it is not uncommon for a large company to use IMS or IDMS and DB2 on the mainframe, Oracle and MySQL on several different UNIX servers, Microsoft SQL Server on Windows servers, as well as pockets of other DBMS products such as Sybase, Ingres, Adabas, and PostgreSQL on various platforms, not to mention single-user PC DBMS products such as Microsoft Access, Paradox, and FileMaker. Who chose to install all these DBMSs and why?

Unfortunately, often the answer is that not much thought and planning went into the decision-making process. Sometimes the decision to purchase and install a new DBMS is driven by a business need or a new application. This is reasonable if your organization has no DBMS and must purchase one for the first time. This is rarely the case, though. Regardless of whether a DBMS exists on-site, a new DBMS is often viewed as a requirement for a new application. Sometimes a new DBMS product is purchased and installed without first examining if the application could be successfully implemented

using an existing DBMS. Or, more likely, the DBAs know the application can be implemented using an existing DBMS but lack the organizational power or support to reject a new DBMS proposal.

There are other reasons for the existence of multiple DBMS platforms in a single organization. Perhaps the company purchased a commercial off-the-shelf application package that does not run on any of the current DBMS platforms. Sometimes the decision to buy a new DBMS is driven by the desire to support the latest and greatest technology. For example, many mainframe shops moving from a hierarchic (IMS) or CODASYL (IDMS) database model to the relational model deployed DB2, resulting in an additional DBMS to learn and support. Then, when client/server computing became popular, additional DBMSs were implemented on UNIX, Linux, and Windows servers.

Once a DBMS is installed, removal can be difficult because of incompatibilities among the different DBMSs and the necessity of converting application code. Furthermore, when a new DBMS is installed, old applications and databases are usually not migrated to it. The old DBMS remains and must continue to be supported. This complicates the DBA's job.

So what should be done? Well, the DBA group should be empowered to make the DBMS decisions for the organization. No business unit should be allowed to purchase a DBMS without the permission of the DBA group. This is a difficult provision to implement and even more difficult to enforce. Business politics often work against the DBA group because it frequently possesses less organizational power than other business executives.

The DBA group should be empowered to make the DBMS decisions for the organization.

Choosing a DBMS

The DBA group should set a policy regarding the DBMS products to be supported within the organization. Whenever possible, the policy should minimize the number of different DBMS products. For a shop with multiple operating systems and multiple types of hardware, choose a default DBMS for the platform. Discourage deviation from the default unless a compelling business case exists—a business case that passes the technical inspection of the DBA group.

Most of the major DBMS products have similar features, and if the feature or functionality does not exist today, it probably will within 18 to 24 months. So, exercise caution before deciding to choose a DBMS based solely on its ability to support a specific feature.

When choosing a DBMS, select a product from a tier-1 vendor.

When choosing a DBMS, it is wise to select a product from a tier-1 vendor as listed in [Table 2.1](#). Tier 1 represents the largest vendors having the most heavily implemented and supported products on the market. You cannot go wrong with DB2 or Oracle. Both are popular and support just about any type of database. Another major player is Microsoft SQL Server, but only for Windows platforms. DB2 and Oracle run on multiple platforms ranging from mainframe to UNIX, as well as Windows and even handheld devices. Choosing a DBMS other than these three should be done only under specific circumstances.

Table 2.1. Tier-1 DBMS Vendors

DBMS Vendor	DBMS Product
IBM Corporation New Orchard Road Armonk, NY 10504 Phone: (914) 499-1900	DB2
Oracle Corporation 500 Oracle Parkway Redwood Shores, CA 94065 Phone: (650) 506-7000	Oracle
Microsoft Corporation One Microsoft Way Redmond, WA 98052 Phone: (425) 882-8080	SQL Server

After the big three come MySQL, Sybase, Teradata, and Informix. [Table 2.2](#) lists these tier-2 DBMS vendors. All of these offerings are quality DBMS products, but their installed base is smaller, their products are engineered and

marketed for niche purposes, or the companies are smaller with fewer resources than the Big Three (IBM, Oracle, and Microsoft), so there is some risk in choosing a DBMS from tier 2 instead of tier 1. However, there may be solid reasons for deploying a tier-2 solution, such as the high performance offered by Informix or the data warehousing and analytics capabilities of Teradata.

Table 2.2. Tier-2 DBMS Vendors

DBMS Vendor	DBMS Product
IBM Corporation New Orchard Road Armonk, NY 10504 Phone: (914) 499-1900	Informix Dynamic Server
Sybase Inc. (an SAP Company) 6475 Christie Avenue Emeryville, CA 94608 Phone: (510) 922-3500	Adaptive Server Enterprise
Teradata Corporation 10000 Innovation Drive Dayton, OH 45342 Phone: (937) 242-4030	Teradata
MySQL (a subsidiary of Oracle Corporation) Phone: (208) 338-8100	MySQL

Of course, there are other DBMS products on the market, many of which are fine products and worthy of consideration for specialty processing, certain predefined needs, and niche roles. If your company is heavily into the open-source software movement, PostgreSQL, EnterpriseDB, or MySQL might be viable options. If an object DBMS is important for a specific project, you might consider ObjectDesign or Versant. And there are a variety of NoSQL DBMS offerings available, too, such as Hadoop, Cassandra, and MongoDB.¹

Choosing any of the lower-tier candidates involves incurring additional risk.

However, for the bulk of your data management needs, a DBMS from a tier-1, or perhaps tier-2, DBMS vendor will deliver sufficient functionality

with minimal risk. A myriad of DBMS products are available, each with certain features that make them worthy of consideration on a case-by-case basis. Choosing any of the lower-tier candidates—even such major names as Software AG’s Adabas and Actian’s Ingres—involves incurring additional risk. Refer to [Appendix B](#) for a list of DBMS vendors.

I do not want it to sound as if the selection of a DBMS is a no-brainer. You *will* need a strategy and a plan for selecting the appropriate DBMS for your specific situation. When choosing a DBMS, be sure to consider each of these factors:

- *Operating system support.* Does the DBMS support the operating systems in use at your organization, including the versions that you are currently using and plan on using?
- *Type of organization.* Take into consideration the corporate philosophy when you choose a DBMS. Some organizations are very conservative and like to keep a tight rein on their environments; these organizations tend to gravitate toward traditional mainframe environments. Government operations, financial institutions, and insurance and health companies usually tend to be conservative. More-liberal organizations are often willing to consider alternative architectures. It is not uncommon for manufacturing companies, dot-coms, and universities to be less conservative. Finally, some companies just do not trust Windows as a mission-critical environment and prefer to use UNIX; this rules out some database vendors (Microsoft SQL Server, in particular).
- *Benchmarks.* What performance benchmarks are available from the DBMS vendor and other users of the DBMS? The Transaction Processing Performance Council (TPC) publishes official database performance benchmarks that can be used as a guideline for the basic overall performance of many different types of database processing. (Refer to the sidebar “[The Transaction Processing Performance Council](#)” for more details.) In general, performance benchmarks can be useful as a broad indicator of database performance but should not be the only determinant when selecting a DBMS. Many of the TPC benchmarks are run against database implementations that are not representative of most production database systems and therefore are not indicative of the actual performance of a particular DBMS. In

addition, benchmarks are constantly updated to show new and improved performance measurements for each of the major DBMS products, rendering the benchmark “winners” obsolete very quickly.

Benchmarks are constantly updated to show new and improved performance measurements.

- *Scalability*. Does the DBMS support the number of users and database sizes you intend to implement? How are large databases built, supported, and maintained—easily or with a lot of pain? Are there independent users who can confirm the DBMS vendor’s scalability claims?
- *Availability of supporting software tools*. Are the supporting tools you require available for the DBMS? These items may include query and analysis tools, data warehousing support tools, database administration tools, backup and recovery tools, performance-monitoring tools, capacity-planning tools, database utilities, and support for various programming languages.

The Transaction Processing Performance Council (TPC)

The Transaction Processing Performance Council is an independent, not-for-profit organization that manages and administers performance benchmark tests. Its mission is to define transaction processing and database benchmarks to provide the industry with objective, verifiable performance data. TPC benchmarks measure and evaluate computer functions and operations.

The definition of *transaction* espoused by the TPC is a business one. A typical TPC transaction includes the database updates for things such as inventory control (goods), airline reservations (services), and banking (money).

The benchmarks produced by the TPC measure performance in terms of how many transactions a given system and database can perform per unit of time, for example, number of transactions per second. The TPC defines three benchmarks:

- TPC-C, for planned production workload in a transaction environment
- TPC-H, a decision support benchmark consisting of a suite of business-

oriented ad hoc queries and concurrent data modifications

- TPC-E, an updated OLTP workload (based on financial transaction processing)

Additional information and in-depth definitions of these benchmarks can be found at the TPC Web site at www.tpc.org (see [Figure 2.1](#)).

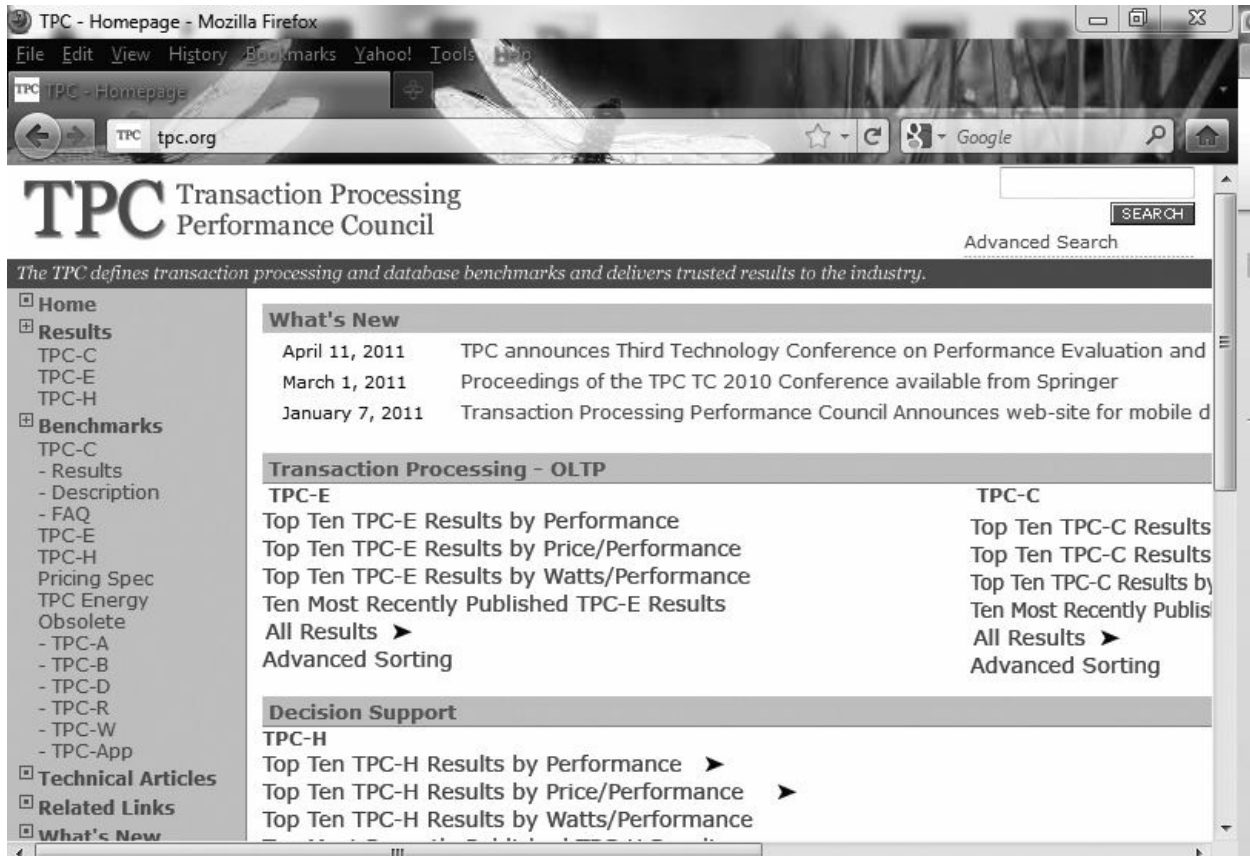


Figure 2.1. The TPC Web site

- *Technicians.* Is there a sufficient supply of skilled database professionals for the DBMS? Consider your needs in terms of DBAs, technical support personnel (system programmers and administrators, operations analysts, etc.), and application programmers.
- *Cost of ownership.* What is the total cost of ownership of the DBMS? DBMS vendors charge wildly varying prices for their technology. Total cost of ownership should be calculated as a combination of the license cost of the DBMS; the license cost of any required supporting software; the cost of database professionals to program, support, and administer

the DBMS; and the cost of the computing resources required to operate the DBMS.

- *Release schedule.* How often does the DBMS vendor release a new version? Some vendors have rapid release cycles, with new releases coming out every 12 to 18 months. This can be good or bad, depending on your approach. If you want cutting-edge features, a rapid release cycle is good. However, if your shop is more conservative, a DBMS that changes frequently can be difficult to support. A rapid release cycle will cause conservative organizations either to upgrade more frequently than they would like or to live with outdated DBMS software that is unlikely to have the same level of support as the latest releases.
- *Reference customers.* Will the DBMS vendor supply current user references? Can you find other users on your own who might provide more impartial answers? Speak with current users to elicit issues and concerns you may have overlooked. How is support? Does the vendor respond well to problems? Do things generally work as advertised? Are there a lot of bug fixes that must be applied continuously? What is the quality of new releases? These questions can be answered only by the folks in the trenches.

When choosing a DBMS, be sure to take into account the complexity of the products. DBMS software is very complex and is getting more complex with each new release. Functionality that used to be supported only with add-on software or independent programs is increasingly being added as features of the DBMS, as shown in [Figure 2.2](#). You will need to plan for and support all the features of the DBMS. Even if there is no current requirement for certain features, once you implement the DBMS the programmers and developers will find a reason to use just about anything the vendor threw into it. It is better to plan and be prepared than to allow features to be used without a plan for supporting them.

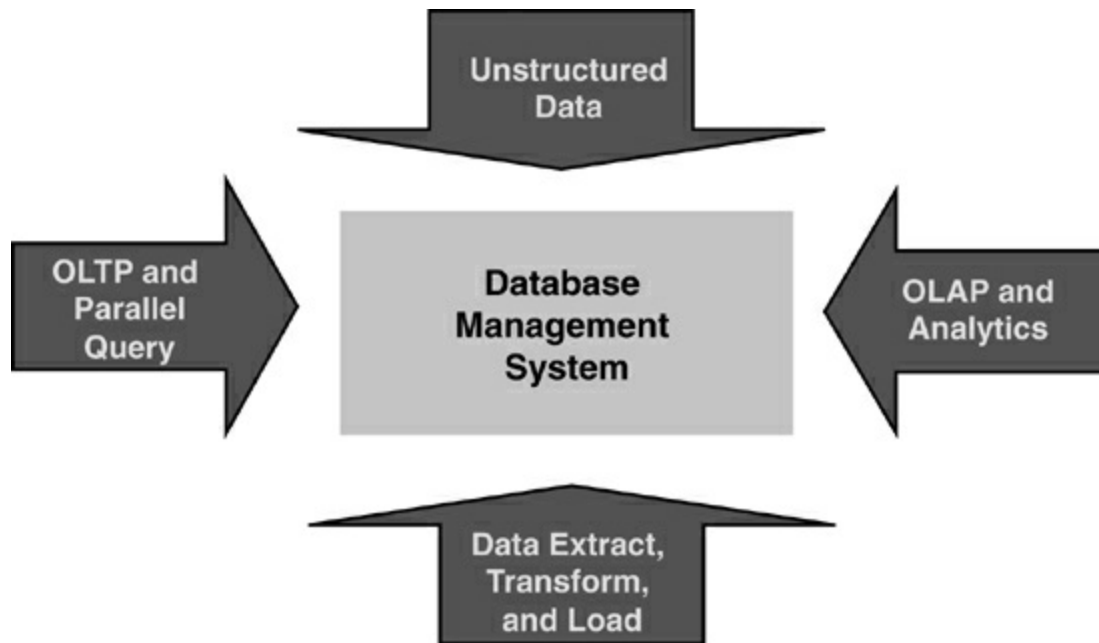


Figure 2.2. Convergence of features and functionality in DBMS software

DBMS Architectures

The supporting architecture for the DBMS environment is very critical to the success of the database applications. One wrong choice or poorly implemented component of the overall architecture can cause poor performance, downtime, or unstable applications.

When mainframes dominated enterprise computing, DBMS architecture was a simpler concern. Everything ran on the mainframe, and that was that. However, today the IT infrastructure is distributed and heterogeneous. The overall architecture—even for a mainframe DBMS—will probably consist of multiple platforms and interoperating system software. A team consisting of business and IT experts, rather than a single person or group, should make the final architecture decision. Business experts should include representatives from various departments, as well as from accounting and legal for software contract issues. Database administration representatives (DA, DBA, and SA), as well as members of the networking group, operating system experts, operations control personnel, programming experts, and any other interested parties, should be included in this team.

The supporting architecture for the DBMS environment is very critical to the success of the database applications.

Furthermore, be sure that the DBMS you select is appropriate for the nature and type of processing you plan to implement. Four levels of DBMS architecture are available: enterprise, departmental, personal, and mobile.

An *enterprise DBMS* is designed for scalability and high performance. An enterprise DBMS must be capable of supporting very large databases, a large number of concurrent users, and multiple types of applications. The enterprise DBMS runs on a large-scale machine, typically a mainframe or a high-end server running UNIX, Linux, or Windows Server. Furthermore, an enterprise DBMS offers all the “bells and whistles” available from the DBMS vendor. Multiprocessor support, support for parallel queries, and other advanced DBMS features are core components of an enterprise DBMS.

Four levels of DBMS architecture are available: enterprise, departmental, personal, and mobile.

A *departmental DBMS*, sometimes referred to as a workgroup DBMS, serves the middle ground. The departmental DBMS supports small to medium-size workgroups within an organization; typically, it runs on a UNIX, Linux, or Windows server. The dividing line between a departmental database server and an enterprise database server is quite gray. Hardware and software upgrades can allow a departmental DBMS to tackle tasks that previously could be performed only by an enterprise DBMS. The steadily falling cost of departmental hardware and software components further contributes to lowering the total cost of operation and enabling a workgroup environment to scale up to serve the enterprise.

A *personal DBMS* is designed for a single user, typically on a low- to medium-powered PC platform. Microsoft Access, SQLite, and FileMaker² are examples of personal database software. Of course, the major DBMS vendors also market personal versions of their higher-powered solutions, such as Oracle Database Personal Edition and DB2 Personal Edition. Sometimes the low cost of a personal DBMS results in a misguided attempt to choose a personal DBMS for a departmental or enterprise solution. However, do not be lured by the low cost. A personal DBMS product is suitable only for very small-scale projects and should never be deployed for multiuser applications.

Finally, the *mobile DBMS* is a specialized version of a departmental or enterprise DBMS. It is designed for remote users who are not usually

connected to the network. The mobile DBMS enables local database access and modification on a laptop or handheld device. Furthermore, the mobile DBMS provides a mechanism for synchronizing remote database changes to a centralized enterprise or departmental database server.

A DBMS designed for one type of processing may be ill suited for other uses. For example, a personal DBMS is not designed for multiple users, and an enterprise DBMS is generally too complex for single users. Be sure to understand the differences among enterprise, departmental, personal, and mobile DBMS software, and choose the appropriate DBMS for your specific data-processing needs. You may need to choose multiple DBMS types—that is, a DBMS for each level—with usage determined by the needs of each development project.

If your organization requires DBMS solutions at different levels, favor the selection of a group of DBMS solutions from the same vendor whenever possible. Doing so will minimize differences in access, development, and administration. For example, favor Oracle Database Personal Edition for your single-user DBMS needs if your organization uses Oracle as the enterprise DBMS of choice.

DBMS Clustering

Clustering is the use of multiple “independent” computing systems working together as a single, highly available system. A modern DBMS offers clustering support to enhance availability and scalability. The two predominant architectures for clustering are *shared-disk* and *shared-nothing*. These names do a good job of describing the nature of the architecture—at least at a high level.

A modern DBMS offers clustering support to enhance availability and scalability.

Shared-nothing clustering is depicted in [Figure 2.3](#). In a shared-nothing architecture, each system has its own private resources (memory, disks, etc.). The clustered processors communicate by passing messages through a network that interconnects the computers. In addition, requests from clients are automatically routed to the system that owns the resource. Only one of the clustered systems can “own” and access a particular resource at a time. In the event a failure occurs, resource ownership can be dynamically transferred

to another system in the cluster. The main advantage of shared-nothing clustering is scalability. In theory, a shared-nothing multiprocessor can scale up to thousands of processors because they do not interfere with one another—nothing is shared.




Figure 2.3. Shared-nothing architecture

The main advantage of shared-nothing clustering is scalability.

Shared-disk clustering is better suited to large-enterprise processing in a mainframe environment.

In a *shared-disk* environment, all the connected systems share the same disk devices, as shown in [Figure 2.4](#). Each processor still has its own private memory, but all the processors can directly address all the disks. Typically, shared-disk clustering does not scale as well for smaller machines as shared-nothing clustering. Shared-disk clustering is better suited to large-enterprise processing in a mainframe environment. Mainframes—very large processors—are capable of processing enormous volumes of work. Great benefits can be obtained with only a few clustered mainframes, while many PC and midrange processors would need to be clustered to achieve similar benefits.




Figure 2.4. Shared-disk architecture

Shared-disk clustering is usually preferable for applications and services requiring only modest shared access to data and for applications or workloads that are very difficult to partition. Applications with heavy data update requirements are probably better implemented as shared-nothing. [Table 2.3](#) compares the capabilities of shared-disk and shared-nothing architectures.

Table 2.3. Comparison of Shared-Disk and Shared-Nothing Architectures

Shared-Disk	Shared-Nothing
Quick adaptability to changing workloads	Can exploit simpler, cheaper hardware
High availability	Almost unlimited scalability
Performs best in a heavy read environment	Works well in a high-volume, read-write environment
Data need not be partitioned	Data is partitioned across the cluster

The major DBMS vendors provide support for different types of clustering with different capabilities and requirements. For example, DB2 for z/OS provides shared-disk clustering with its Data Sharing and Parallel Sysplex capabilities; DB2 on non-mainframe platforms uses shared-nothing clustering. Oracle's Real Application Clusters provide shared-disk clustering.

For most users, the primary benefit of clustering is the enhanced availability that accrues by combining processors. In some cases, clustering can help an enterprise to achieve five-nines (99.999 percent) availability. Additionally, clustering can be used for load balancing and failover.

DBMS Proliferation

As a rule of thumb, create a policy (or at least some simple guidelines) that must be followed before a new DBMS can be brought into the organization. Failure to do so can cause a proliferation of different DBMS products that will be difficult to support. It can also cause confusion regarding which DBMS to use for which development effort.

A proliferation of different DBMS products can be difficult to support.

As mentioned earlier, there is a plethora of DBMS vendors, each touting its benefits. As a DBA, you will be bombarded with marketing and sales efforts that attempt to convince you that you need another DBMS. Try to resist unless a very compelling reason is given and a short-term return on investment (ROI) can be demonstrated. Even when confronted with valid reasons and good ROI, be sure to double-check the arguments and ROI calculations. Sometimes the reasons specified are outdated and the ROI figures do not take everything into account—such as the additional cost of administration.

Remember, every DBMS requires database administration support. Moreover, each DBMS uses different methods to perform similar tasks. The fewer DBMS products installed, the less complicated database administration becomes, and the better your chances become of providing effective data management resources for your organization.

Hardware Issues

When establishing a database environment for application development, selecting the DBMS is only part of the equation. The hardware and operating system on which the DBMS will run will greatly impact the reliability, availability, and scalability (RAS) of the database environment. For example, a mainframe platform such as an IBM zEC12 running z/OS will probably provide higher RAS than a midrange IBM xSeries machine running AIX, which in turn will probably exceed a Dell server running Windows. That is not to say everything should run on a mainframe; other issues such as cost, experience, manageability, and the needs of the applications to be developed must be considered. The bottom line is that you must be sure to factor hardware platform and operating system constraints into the DBMS selection

criteria.

Factor hardware platform and operating system constraints into the DBMS selection criteria.

Cloud Database Systems

Cloud computing (see the sidebar) is increasing in usage, especially at small to medium-size businesses. A cloud implementation can be more cost-effective than building an entire local computing infrastructure that requires management and support.

A cloud database system delivers DBMS services over the Internet. The trade-off essentially comes down to trusting a cloud provider to store and manage your data in return for minimizing database administration and maintenance cost and effort. Using cloud database systems can enable organizations, especially smaller ones without the resources to invest in an enterprise computing infrastructure, to focus on their business instead of their computing environment.

By consolidating data sources in the cloud, it is possible to improve collaboration among partners, branch offices, remote workers, and mobile devices, because the data becomes accessible as a service. There is no need to install, set up, patch, or manage the DBMS software because the cloud provider manages and cares for these administrative tasks. Of course, the downside is that your data is now stored and controlled by an external agent—the cloud provider. Another inherent risk of cloud computing is the possibility of nefarious agents posing as legitimate customers.

Cloud Computing Overview

At a high level, cloud computing is the delivery of computing as a service. Cloud computing applications rely on a network (typically the Internet) to provide users with shared resources, software, and data. With cloud computing, computer systems and applications are supposed to function like a utility provider (such as the electricity grid).

The term *cloud* is used as a metaphor for the Internet. It is based on the tendency to draw network access as an abstract “cloud” in infrastructure diagrams. An example of this can be seen in [Figure 1.11](#) in [Chapter 1](#) of this book.

From a DBMS perspective, cloud computing moves the data and its management away from your local computing environment and delivers it as a service over the Internet.

An example of a cloud database platform is Microsoft SQL Azure. It is built on SQL Server technologies and is a component of the Windows Azure platform.

Installing the DBMS

Once the DBMS has been chosen, you will need to install it. Installing a DBMS is not as simple as popping a CD into a drive and letting the software install itself (or, for you mainframe folks, just using IEBGENER to copy it from a tape). A DBMS is a complex piece of software that requires up-front planning for installation to be successful. You will need to understand the DBMS requirements and prepare the environment for the new DBMS.

DBMS Installation Basics

The very first thing to do when you install a DBMS for the first time is to understand the prerequisites. Every DBMS comes with an installation manual or guide containing a list of the operating requirements that must be met for the DBMS to function properly. Examples of prerequisites include ensuring that an appropriate version of the operating system is being used, verifying that there is sufficient memory to support the DBMS, and ensuring that any related software to be used with the DBMS is the proper version and maintenance level.

Read the installation guide from cover to cover.

Once the basics are covered, read the installation guide from cover to cover. Make sure that you understand the process before you even begin to install the DBMS. Quite a few preparations need to be made before installing a DBMS, and reading about them *before* you start will ensure a successful installation. Review how the installation program or routine for the DBMS operates, and follow the explicit instructions in the installation guide provided with the DBMS software. You additionally might want to work closely with the DBMS vendor during an initial installation to ensure that your plans are sound. In some cases, working with a local, experienced

vendor or consultant can be beneficial to avoid installation and configuration errors.

The remainder of this section will discuss some of the common preparations that are required before a DBMS can be installed. If the DBMS is already operational and you are planning to migrate to a new DBMS release, refer to the section “[Upgrading DBMS Versions and Releases.](#)”

Hardware Requirements

Every DBMS has a basic CPU requirement, meaning a CPU version and minimum processor speed required for the DBMS to operate. Additionally, some DBMSs specify hardware models that are required or unsupported. Usually the CPU criterion will suffice for an Intel environment, but in a mainframe or enterprise server environment the machine model can make a difference with regard to the DBMS features supported. For example, certain machines have built-in firmware that can be exploited by the DBMS if the firmware is available.

Choose the correct DBMS for your needs and match your hardware to the requirements of the DBMS.

Furthermore, each DBMS offers different “flavors” of its software for specific needs. (I use “flavor” as opposed to “version” or “release,” which specify different iterations of the same DBMS.) Different flavors of the DBMS (at the same release level) are available for specific environments such as parallel processing, pervasive computing (such as handheld devices), data warehousing, and/or mobile computing. Be sure to choose the correct DBMS for your needs and to match your hardware to the requirements of the DBMS.

Storage Requirements

A DBMS requires disk storage to run. And not just for the obvious reason—to create databases that store data. A DBMS will use disk storage for the indexes to be defined on the databases as well as for the following items:

- The system catalog or data dictionary used by the DBMS to manage and track databases and related information. The more database objects you plan to create, the larger the amount of storage required by the system catalog.

- Any other system databases required by the DBMS, for example, to support distributed connections or management tools.
- Log files that record all changes made to every database. These include active logs, archive logs, rollback segments, and any other type of change log required by the DBMS.
- Start-up or control files that must be accessed by the DBMS when it is started or initialized.
- Work files used by the DBMS to sort data or for other processing needs.
- Default databases used by the DBMS for system structures or as a default catchall for new database objects as they are created.
- Temporary database structures used by the DBMS (or by applications accessing databases) for transient data that is not required to be persistent but needs reserved storage during operations (such as rebuilding clustered indexes on Microsoft SQL Server).
- System dump and error-processing files.
- DBA databases used for administration, monitoring, and tuning—for example, DBA databases used for testing new releases, migration scripts, and so on.

Be sure to factor in every storage requirement of the DBMS and reserve the appropriate storage. Also, be aware that the DBMS will use many of these databases and file structures concurrently. Therefore, it is a good idea to plan on using multiple storage devices even if you will not fill them to capacity. Proper database and file placement will enable the DBMS to operate more efficiently because concurrent activities will not be constrained by the physical disk as data is accessed.

Factor in every storage requirement of the DBMS and reserve the appropriate storage.

Disk storage is not the only requirement of a DBMS. Tape or optical discs (such as DVDs and CDs) are also required for tasks such as database backups and log off-loading. When the active log file fills up, the log records must be off-loaded to an archive log either on disk or on tape, as shown in [Figure 2.5](#). Depending on the DBMS being used and the features that have been activated, this process may be automatic or manual. The archive log files

must be retained for recovery purposes, and even if originally stored on disk, they must eventually be migrated to an external storage mechanism for safekeeping.

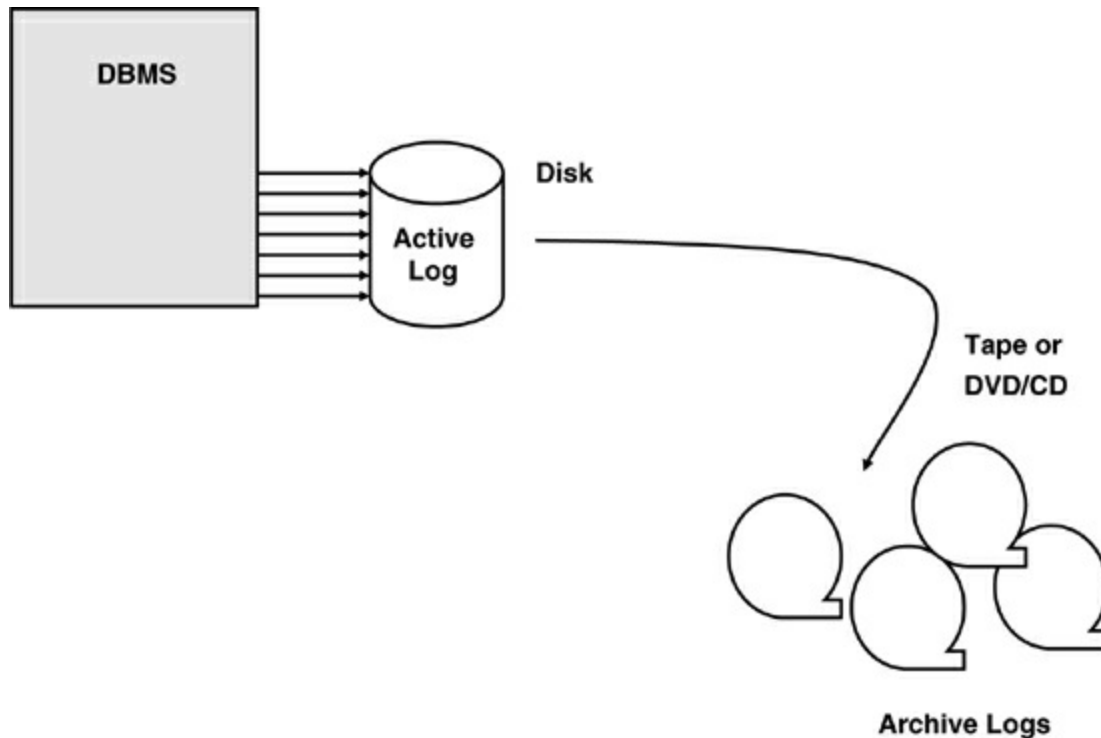


Figure 2.5. Log off-loading

Plan on maintaining multiple tape or CD/DVD drives to enable the DBMS to run concurrent multiple processes that require external storage, such as concurrent database backups. Database outages can occur if you single-thread your database backup jobs using a single drive.

Memory Requirements

Relational DBMSs, as well as their databases and applications, love memory. A DBMS requires memory for basic functionality and will use it for most internal processes such as maintaining the system global area and performing many DBMS tasks.

A DBMS requires a significant amount of memory to cache data in memory structures in order to avoid I/O. Reading data from a disk storage device is always more expensive and slower than moving the data around in memory. [Figure 2.6](#) shows how the DBMS uses a memory structure called a [buffer pool](#) or [data cache](#) to reduce physical I/O requests. By caching data

that is read into a buffer pool, the DBMS can avoid I/O for subsequent requests for the same data, as long as it remains in the buffer pool. In general, the larger the buffer pool, the longer the data can remain in memory and the better overall database processing will perform.

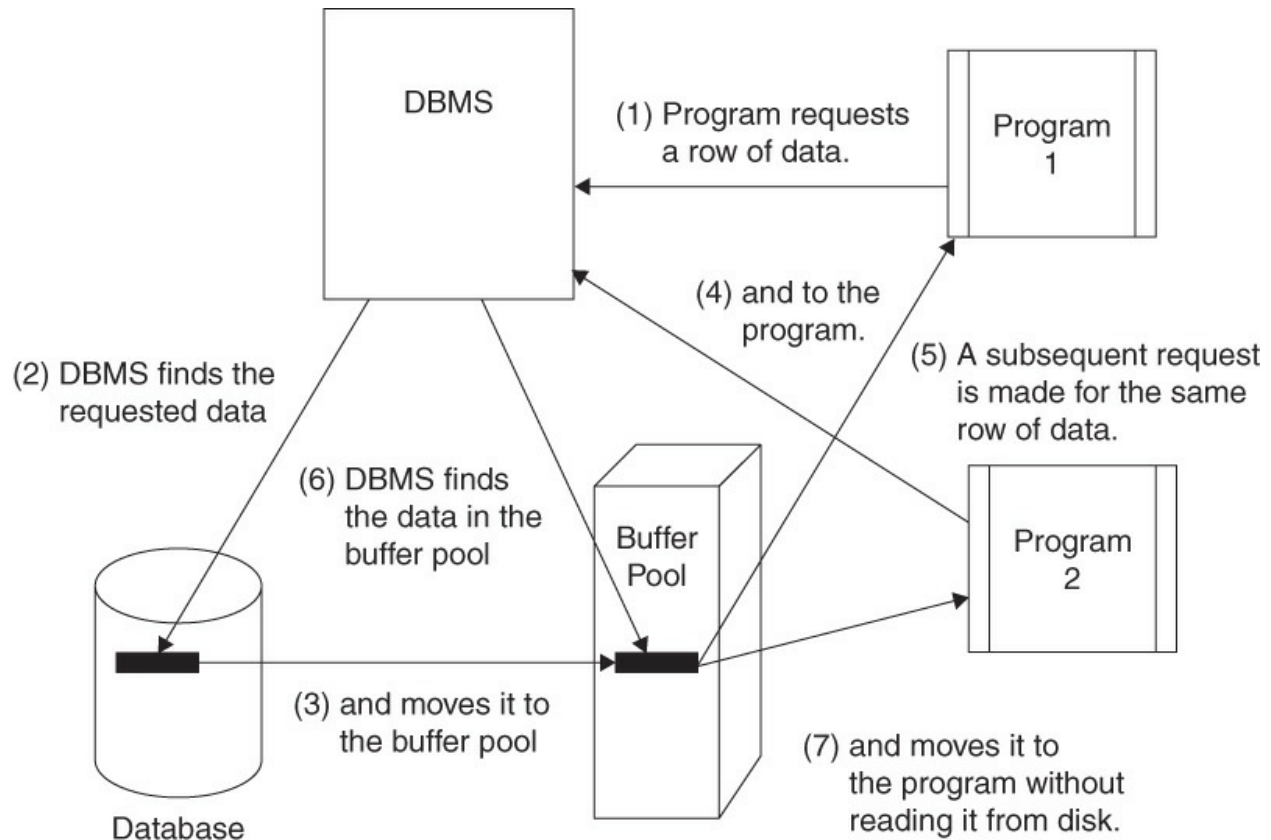


Figure 2.6. Buffer pool (or data cache)

Besides data, the DBMS will cache other structures in memory. Most DBMSs set aside memory to store program structures required by the DBMS to process database requests.³ The *program cache* stores things like “compiled” SQL statements, database authorizations, and database structure blocks that are used by programs as they are executed. When these structures are cached, database processing can be optimized because additional I/O requests to access them from a physical storage device are avoided.

Memory is typically required by the DBMS to support other features such as handling lock requests, facilitating distributed data requests, sorting data, optimizing processes, and processing SQL.

Ensure that the DBMS has a more-than-adequate supply of memory at its disposal. This will help to optimize database processing and minimize

potential problems.

Ensure that the DBMS has a more-than-adequate supply of memory at its disposal.

Configuring the DBMS

Configuring the system parameters of the DBMS controls the manner in which the DBMS functions and the resources made available to it.⁴ Each DBMS allows its system parameters to be modified in different ways, but the installation process usually sets the DBMS system parameters by means of radio buttons, menus, or panel selections. During the installation process, the input provided to the installation script will be used to establish the initial settings of the system parameters.

Each DBMS also provides a method to change the system parameters once the DBMS is operational.

Each DBMS also provides a method to change the system parameters once the DBMS is operational. Sometimes you can use DBMS commands to set the system's parameters; sometimes you must edit a file that contains the current system parameter settings. If you must edit a file, be very careful: An erroneous system parameter setting can be fatal to the operational status of the DBMS.

What do the system parameters control? Well, for example, system parameters control DBA authorization to the DBMS and the number of active database logs; system parameters set the amount of memory used for data and program caching and turn DBMS features on or off. Although every DBMS has system parameters that control its functionality, each DBMS has a different method of setting and changing the values. And, indeed, each DBMS has different specifications that can be set using system parameters.

Beware of simply using default system parameters when installing the database system software. Although using defaults can save time and make for an easier installation, it can also result in subsequent problems. Most DBMSs are poorly served, in the long run, by default settings and, in some cases, can experience worsening performance over time because resources were not preallocated during installation or setup.

Be sure to understand fully the parameters used by your DBMS. Failure to

do so can result in an incorrectly configured database environment, which can cause performance problems, data integrity problems, or even DBMS failure.

Connecting the DBMS to Supporting Infrastructure Software

Part of the DBMS installation process is the connection of the DBMS to other system software components that must interact with the DBMS. Typical infrastructure software that may need to be configured to work with the DBMS includes networks, transaction processing monitors, message queues, other types of middleware, programming languages, systems management software, operations and job control software, Web servers, and application servers.

Each piece of supporting infrastructure software will have different requirements for interfacing with the DBMS.

Each piece of supporting infrastructure software will have different requirements for interfacing with the DBMS. Typical configuration procedures can include installing DLL files, creating new parameter files to establish connections, and possibly revisiting the installation procedures for the supporting software to install components required to interact with the DBMS.

Installation Verification

After installing the DBMS, you should run a battery of tests to verify that the DBMS has been properly installed and configured. Most DBMS vendors supply sample programs and installation verification procedures for this purpose. Additionally, you can ensure proper installation by testing the standard interfaces to the DBMS. One standard interface supported by most DBMSs is an interactive SQL interface where you can submit SQL statements directly to the DBMS.⁵

Create a set of SQL code that comprises SELECT, INSERT, UPDATE, and DELETE statements issued against sample databases. Running such a script after installation helps you to verify that the DBMS is installed correctly and operating as expected.

Furthermore, be sure to verify that all required connections to supporting software are operational and functioning properly. If the DBMS vendor does

not supply sample programs, you may need to create and run simple test programs for each environment to ensure that the supporting software connections are functioning correctly with the DBMS.

DBMS Environments

Generally, installing a DBMS involves more than simply installing one instance or subsystem. To support database development, the DBA needs to create multiple DBMS environments to support, for example, testing, quality assurance, integration, and production work. Of course, it is possible to support multiple environments in a single DBMS instance, but it is not prudent. Multiple DBMS installations are preferable to support multiple development environments for a single database. This minimizes migration issues and won't require complex database naming conventions to support. Furthermore, segregating database instances makes testing, tuning, and monitoring easier.

Upgrading DBMS Versions and Releases

Change is a fact of life, and each of the major DBMS products changes quite rapidly. A typical release cycle for DBMS software is 18 to 24 months for major releases, with constant bug fixes and maintenance updates delivered between major releases. Indeed, keeping DBMS software up-to-date can be a full-time job.

Change is a fact of life.

The DBA must develop an approach to upgrading DBMS software that conforms to the organization's needs and minimizes business disruptions due to outages and database unavailability.

You may have noticed that I use the terms [version](#) and [release](#) somewhat interchangeably. That is fine for a broad discussion of DBMS upgrades, but a more precise definition is warranted. For a better discussion of the differences between a version and a release, please refer to the sidebar.

A DBMS version upgrade can be thought of as a special case of a new installation. All the procedures required of a new installation apply to an upgrade: You must plan for appropriate resources, reconsider all system parameters, and ensure that all supporting software is appropriately connected. However, another serious issue must be planned for: existing

users and applications. An upgrade needs to be planned to cause as little disruption to the existing users as possible. Furthermore, any additional software that works with the DBMS (such as purchased applications, DBA tools, utilities, and so on) must be verified to be compatible with the new DBMS version. Therefore, upgrading can be a tricky and difficult task.

Version or Release?

Vendors typically make a distinction between a version and a release of a software product. A new version of software is a major concern, with many changes and new features. A release is typically minor, with fewer changes and not as many new features.

For example, moving from Version 10g of Oracle Database to Version 11g would be a major change—a version change. However, an in-between point such as Oracle Database 11g Release 2 would be considered a release—consisting of a smaller number of changes. Usually DBMS vendors increase prices for versions, but not necessarily for releases (but that is not a hard-and-fast rule).

Usually significant functionality is added for version upgrades, less so for point releases. Nevertheless, upgrading from one point release to another can have just as many potential pitfalls as a version upgrade. It depends on the nature of the new features provided in each specific release.

The issues and concerns discussed in this chapter pertain to both types of DBMS upgrades: to a new release and to a new version.

In a complex, heterogeneous, distributed database environment, a coherent upgrade strategy is essential. Truthfully, even organizations with only a single DBMS should approach DBMS upgrades cautiously and plan accordingly. Failure to plan a DBMS upgrade can result in improper and inefficient adoption of new features, performance degradation of new and existing applications, and downtime.

Upgrading to a new DBMS release offers both rewards and risks. The following are some of the benefits of moving to a new release:

- Developers can avail themselves of new features and functionality delivered only in the new release. If development requires a new

feature, or can simply benefit from a new feature, program development time can be reduced or made more cost-effective.

- For purchased applications, the application vendor may require a specific DBMS version or release for specific versions of its application to enable specific functionality within the application.
- New DBMS releases usually deliver enhanced performance and availability features that can optimize existing applications. Sometimes a new DBMS release is required to scale applications to support additional users or larger amounts of data.
- DBMS vendors often provide better support and respond to problems faster for a new release of their software. DBMS vendors are loath to allow bad publicity about bugs in a new and heavily promoted version of their products.
- Cost savings may accrue by upgrading to a new DBMS release. Some vendors charge additionally when a company uses multiple versions of a DBMS, such as the new version in a test environment and the old in production. When both are migrated to the same version, the price tag for the DBMS sometimes can be reduced.
- Production migration to a new DBMS release will align the test and production database environments, thereby providing a consistent environment for development and implementation. If a new release is running in the test environment for too long, database administration and application development tasks become more difficult because the test databases will operate differently from the production databases.

However, an effective DBMS upgrade strategy must balance the benefits against the risks of upgrading to arrive at the best timeline for migrating to a new DBMS version or release. The risks of upgrading to a new DBMS release include the following:

An effective DBMS upgrade strategy must balance the benefits against the risks of upgrading.

- An upgrade to the DBMS usually involves some level of disruption to business operations. At a minimum, databases will not be available while the DBMS is being upgraded. This can result in downtime and lost business opportunities if the DBMS upgrade occurs during normal

business hours (or if there is no planned downtime). Clustered database implementations may permit some database availability while individual database clusters are migrated to the new DBMS version.

- Other disruptions can occur, such as having to convert database structures or discovering that previously supported features were removed from the new release (thereby causing application errors). Delays to application implementation timelines are another possibility.
- The cost of an upgrade can be a significant barrier to DBMS release migration. First, the cost of the new version or release must be budgeted for (price increases for a new DBMS version can amount to as much as 10 to 25 percent). The upgrade cost must also factor in the costs of planning, installing, testing, and deploying not just the DBMS but also any applications that use databases. Finally, be sure to include the cost of any new resources (such as memory, storage, additional CPUs) required to use the new features delivered by the new DBMS version.⁶
- DBMS vendors usually tout the performance gains that can be achieved with a new release. However, when SQL optimization techniques change, it is possible that a new DBMS release will generate SQL access paths that perform worse than before. DBAs must implement a rigorous testing process to ensure that new access paths are helping, not harming, application performance. When performance suffers, application code may need to be changed—a very costly and time-consuming endeavor. A rigorous test process should be able to catch most of the access path changes in the test environment.
- New DBMS releases may cause features and syntax that are being used in existing applications to be deprecated.⁷ When this occurs, the applications must be modified before migration to the new release can proceed.
- To take advantage of improvements implemented in a new DBMS release, the DBA may have to apply some invasive changes. For example, if the new version increases the maximum size for a database object, the DBA may have to drop and recreate that object to take advantage of the new maximum. This will be the case when the DBMS adds internal control structures to facilitate such changes.
- Supporting software products may lack immediate support for a new

DBMS release. Supporting software includes the operating system, transaction processors, message queues, purchased applications, DBA tools, development tools, and query and reporting software.

After weighing the benefits of upgrading against the risks of a new DBMS release, the DBA group must create an upgrade plan that works for the organization. Sometimes the decision will be to upgrade immediately upon availability, but often there is a lag between the general availability of a new release and its widespread adoption.

When the risks of a new release outweigh the benefits, some organizations may decide to skip an interim release if doing so does not impact a future upgrade. For example, a good number of Oracle customers migrated directly from Oracle7 to Oracle8i, skipping Oracle8. If the DBMS vendor does not allow users to bypass a version or release, it is still possible to “skip” a release by waiting to implement that release until the next release is available. For example, consider the following scenario:

1. ABC Corporation is using DB Version 8 from DBCorp.
2. DBCorp announces Version 9 of DB.
3. ABC Corporation analyzes the features and risks and determines not to upgrade immediately.
4. DBCorp later announces DB Version 10 and that no direct migration path will be provided from Version 8 to Version 10.
5. ABC Corporation decides that DB Version 10 provides many useful features and wants to upgrade its current Version 8 implementation of DB. However, it has no compelling reason to first implement and use Version 9.
6. To fulfill its requirements, ABC Corporation first upgrades Version 8 to Version 9 and then immediately upgrades Version 9 to Version 10.

Although a multiple-release upgrade takes more time, it allows customers to effectively control when and how they will migrate to new releases of a DBMS instead of being held hostage by the DBMS vendor. When attempting a multiple-release upgrade of this type, be sure to fully understand the features and functionality added by the DBMS vendor for each interim release. In the case of the hypothetical ABC Corporation, the DBAs would need to research and prepare for the new features of not just Version 10 but also Version 9.

A multiple-release upgrade allows customers to effectively control when and how they will migrate to new releases of a DBMS.

An appropriate DBMS upgrade strategy depends on many things. The following sections outline the issues that must be factored into an effective DBMS release upgrade strategy.

Features and Complexity

Perhaps the biggest factor in determining when and how to upgrade to a new DBMS release is the functionality supported by the new release. Tightly coupled to functionality is the inherent complexity involved in supporting and administering new features.

It is more difficult to delay an upgrade if application developers are clamoring for new DBMS features. If DBMS functionality can minimize the cost and effort of application development, the DBA group will feel pressure to migrate swiftly to the new release. An additional factor that will coerce rapid adoption of a new release is when DBMS problems are fixed in the new release (instead of through regular maintenance fixes).

Regardless of a new release's "bells and whistles," certain administration and implementation details must be addressed before upgrading. The DBA group must ensure that standards are modified to include the new features, educate developers and users as to how new features work and should be used, and prepare the infrastructure to support the new DBMS functionality.

The types of changes required to support the new functionality must be factored into the upgrade strategy. When the DBMS vendor makes changes to internal structures, data page layouts, or address spaces, the risks of upgrading are greater. Additional testing is warranted in these situations to ensure that database utilities, DBA tools, and data extraction and movement tools still work with the revised internal structures.

Complexity of the DBMS Environment

The more complex your database environment is, the more difficult it will be to upgrade to a new DBMS release. The first complexity issue is the size of the environment. The greater the number of database servers, instances, applications, and users, the greater the complexity. Additional concerns include the types of applications being supported. A DBMS upgrade is easier

to implement if only simple, batch-oriented applications are involved. As the complexity and availability requirements of the applications increase, the difficulty of upgrading also increases.

Location of the database servers also affects the release upgrade strategy. Effectively planning and deploying a DBMS upgrade across multiple database servers at various locations supporting different lines of business is difficult. It is likely that an upgrade strategy will involve periods of supporting multiple versions of the DBMS at different locations and for different applications. Supporting different versions in production should be avoided, but that is not always possible.

Finally, the complexity of the applications that access your databases must be considered. The more complex your applications are, the more difficult it will be to ensure their continuing uninterrupted functionality when the DBMS is modified. Complexity issues include the following:

- Usage of stored procedures and user-defined functions.
- Complexity of the SQL—the more tables involved in the SQL and the more complex the SQL features, the more difficult it becomes to ensure that access path changes do not impact performance.
- Client/server processing—network usage and usage of multiple tiers complicates testing the new DBMS release.
- Applications that are designed, coded, and generated by a framework or an IDE (for example, Hibernate) may have additional components that need to be tested with a new DBMS release.
- Integration with other infrastructure software such as message queues and transaction processors can complicate migration because new versions of these products may be required to support the new DBMS release.
- The language used by the programs might also impact DBMS release migration due to different support for compiler versions, changes to APIs (application programming interfaces), or new ways of embedding SQL into application programs.

Reputation of the DBMS Vendor

DBMS vendors have different reputations for technical support, fixing bugs, and responding to problems, which is why customer references are so

important when choosing a database.

The better the reputation of the vendor, the greater the likelihood of organizations rapidly adopting a new release.

The better the reputation of the vendor, the greater the likelihood of organizations rapidly adopting a new release. If the DBMS vendor is good at responding to problems and supporting its customers as they migrate to new releases, those customers will more actively engage in migration activities.

Support Policies of the DBMS

As new releases are introduced, DBMS vendors will retire older releases and no longer support them. The length of time that the DBMS vendor will support an old release must be factored into the DBMS release migration strategy. You should never run a DBMS release in production that is no longer supported by the vendor. If problems occur, the DBMS vendor will not be able to resolve them for you.

Sometimes a DBMS vendor will provide support for a retired release on a special basis and at an increased maintenance charge. If you absolutely must continue using a retired DBMS release (for business or application issues), be sure to investigate the DBMS vendor's policies regarding support for retired releases of its software.

Organization Style

Every organization displays characteristics that reveal its style when it comes to adopting new products and technologies. Industry analysts at Gartner, Inc., have ranked organizations into three distinct groups labeled types A, B, and C. A type-A enterprise is technology driven and, as such, is more likely to risk using new and unproven technologies to try to gain a competitive advantage. A type-B organization is less willing to take risks but will adopt new technologies once others have shaken out the bugs. Finally, a type-C enterprise, very conscious of cost and averse to risk, will lag behind the majority when it comes to migrating to new technology.

Only type-A organizations should plan on moving aggressively to new DBMS releases immediately upon availability and only if the new features of the release will deliver advantages to the company. Type-C enterprises should adopt a very conservative strategy to ensure that the DBMS release is

stable and well tested by type-A and type-B companies first. Type-B organizations will fall somewhere between types A and C: Almost never upgrading immediately, the type-B company will adopt the new release after the earliest users have shaken out the biggest problems, but well before type-C enterprises.

DBA Staff Skill Set

Upgrading the DBMS is easier if your DBA staff is highly skilled and/or experienced. The risk of an upgrade increases as the skills of the DBA staff decrease. If your DBAs are not highly skilled, or have never migrated a DBMS to a new release, consider augmenting your DBA staff with consultants for the upgrade. Deploying an integrated team of internal DBAs and consultants will ensure that your upgrade goes as smoothly as possible. Furthermore, the DBA staff will be better prepared to handle the future upgrades alone.

The risk of an upgrade increases as the skills of the DBA staff decrease.

If consultants will be required, be sure to include their contracting cost in the DBMS release upgrade budget. The budget should allow you to retain the consultants until all production database environments are stable.

Platform Support

When a DBMS vendor unleashes a new release of its product, not all platforms and operating systems are immediately supported. The DBMS vendor usually first supports the platforms and operating systems for which it has the most licensed customers. The order in which platforms are supported for a new release is likely to differ for each DBMS vendor. For example, Linux for System z is more strategic to IBM than to Oracle, so a new DB2 release will most likely support Linux for System z very quickly, whereas this may not be true of Oracle. The issue is even thornier for UNIX platforms because of the sheer number of UNIX variants in the marketplace. The most popular variants are Oracle's Solaris, IBM's AIX, Hewlett-Packard's HP-UX, and Linux, the open-source version of UNIX (the Red Hat and Suse distributions are supported more frequently and rapidly than others). Most DBMS vendors will support these UNIX platforms quickly upon general availability. Other less popular varieties of UNIX will take longer for the

DBMS vendors to support.

When planning your DBMS upgrade, be sure to consider the DBMS platforms you use and try to gauge the priority of your platform to your vendor. Be sure to build some lag time into your release migration strategy to accommodate the vendor's delivery schedule for your specific platforms.

Supporting Software

Carefully consider the impact of a DBMS upgrade on any supporting software. Supporting software includes purchased applications, DBA tools, reporting and analysis tools, and query tools. Each software vendor will have a different time frame for supporting and exploiting a new DBMS release. Review the sidebar to understand the difference between support and exploitation of a new DBMS release.

Carefully consider the impact of a DBMS upgrade on any supporting software.

Some third-party tool vendors follow guidelines for supporting and exploiting new DBMS releases. Whenever possible, ask your vendors to state their policies for DBMS upgrade support. Your vendors will probably not commit to any firm date or date range to support new versions and releases—some DBMS versions are larger and more complicated and therefore take longer to fully exploit.

Support versus Exploit

Some vendors differentiate specifically between supporting and exploiting a new DBMS version or release. Software that supports a new release will continue to function the same as before the DBMS was upgraded, but with no new capabilities. Therefore, if a DBA tool, for example, *supports* a new version of Oracle, it can provide all the services it did for the last release, as long as none of the new features of the new version of Oracle are used. In contrast, a DBA tool that *exploits* a new version or release provides the requisite functionality to operate on the new features of the new DBMS release.

So, to use a concrete example, IBM added support for hashing in Version 10 of DB2. A DBA tool can *support* DB2 Version 10 without operating on hashes, but it must operate on hashes to *exploit* DB2

Version 10.

Prior to migrating to a new DBMS version or release, make sure you understand the difference between supporting and exploiting a new version, and get a schedule for both from your third-party vendors for the DBA tools you use.

Fallback Planning

Each new DBMS version or release should come with a manual that outlines the new features of the release and describes the fallback procedures to return to a prior release of the DBMS. Be sure to review the fallback procedures provided by the DBMS vendor in its release guide. You may need to return to a previous DBMS release if the upgrade contains a bug, performance problems ensue, or other problems arise during or immediately after migration. Keep in mind that fallback is not always an option for every new DBMS release.

If fallback is possible, follow the DBMS vendor's recommended procedures to enable it. You may need to delay the implementation of certain new features for fallback to remain an option. Understand fully the limitations imposed by the DBMS vendor on fallback, and exploit new features only when fallback is no longer an option for your organization.

Migration Verification

The DBA should implement procedures—similar to those for a new installation—to verify that the DBMS release upgrade is satisfactory. Perform the same steps as with a brand-new DBMS install, but also test a representative sampling of your in-house applications to verify that the DBMS upgrade is working correctly and performing satisfactorily.

The DBA should implement procedures to verify that the DBMS release upgrade is satisfactory.

The DBMS Upgrade Strategy

In general, design your DBMS release upgrade policy according to the guidelines discussed in the preceding sections. Each specific DBMS upgrade will be unique, but the strategies we've discussed will help you to achieve success more readily. A well-thought-out DBMS upgrade strategy will

prepare you to support new DBMS releases with minimum impact on your organization and in a style best suited to your company.

Database Standards and Procedures

Before a newly installed DBMS can be used effectively, standards and procedures must be developed for database usage. Studies have shown that companies with high levels of standardization reduce the cost of supporting end users by as much as 35 percent or more as compared to companies with low levels of standardization.

Standards and procedures must be developed for database usage.

Standards are common practices that ensure the consistency and effectiveness of the database environment, such as database naming conventions. *Procedures* are defined, step-by-step instructions that direct the processes required for handling specific events, such as a disaster recovery plan. Failure to implement database standards and procedures will result in a database environment that is confusing and difficult to manage.

The DBA should develop database standards and procedures as a component of corporate-wide IT standards and procedures. They should be stored together in a central location as a printed document, in an online format, or as both. Several vendors offer “canned” standards and procedures that can be purchased for specific DBMS products.

Database Naming Conventions

One of the first standards to be implemented should be a set of guidelines for the naming of database objects. Without standard database object naming conventions, it will be difficult to identify database objects correctly and to perform the proper administration tasks.

Database object naming standards should be developed in conjunction with all other IT naming standards in your organization. In all cases, database naming standards should be developed in cooperation with the data administration department (if one exists) and, wherever possible, should peacefully coexist with other IT standards, but not at the expense of impairing the database environment. For example, many organizations have shop conventions for naming files, but coordinating the database object to the

operating system file may require a specific format for database filenames that does not conform to the shop standards (see [Figure 2.7](#)). Therefore, it may be necessary to make exceptions to existing shop standards for naming database files.

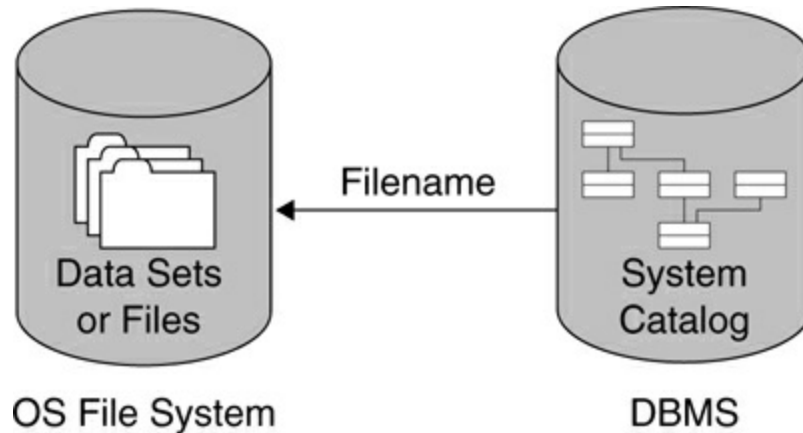


Figure 2.7. Database objects map to filenames

Be sure to create and publish naming standards for all database objects that can be created within each DBMS used by your organization. A basic list of database objects supported by most DBMSs includes databases, tables, columns, views, indexes, constraints, programs, user-defined data types, user-defined functions, triggers, and stored procedures. However, this list is incomplete because each DBMS uses other database objects specific to its operation. For example, DB2 uses plans and storage groups; Oracle uses database links and clusters; SQL Server uses filegroups and rules (see the sidebar).

Be sure to establish naming conventions for all database objects.

The database naming standard should be designed to minimize name changes across environments. For example, embedding a *T* into the name for “test” and a *P* for “production” is a bad idea. It is especially important to avoid this approach for user-visible database objects such as columns, tables, and views. Minimizing name changes simplifies the migration of databases from one environment to another. It is possible to make all the database object names the same by assigning each environment to a different instance or subsystem. The instance or subsystem name, rather than the database object names, will differentiate the environments.

Minimize name changes across environments.

Example Nonstandard Database Objects

Unless you use all three of DB2, Oracle, and SQL Server, some of the database objects that are specific to only one of these database systems probably will be unfamiliar to you. Given that, this sidebar offers short definitions of the database objects mentioned in this section.

For DB2:

- A **plan** is associated with a DB2 application program and refers to packages that contain bound access path details for the SQL in that program.
- A **storage group** is a database object used to associate disk storage with DB2 tablespaces.

For Oracle:

- A **database link** is a schema object in one database that enables you to access objects in another database.
- A **cluster** is made up of a group of tables that share the same data blocks. The tables are grouped together because they share common columns and are often used together.

For SQL Server:

- Database objects and files can be grouped together in **filegroups** for allocation and administration purposes.
 - A **rule** is a freestanding database constraint that can be attached to columns. Microsoft has indicated that rules will be removed from a future version of SQL Server.
-

In most cases, for objects not accessed by typical end users, provide a way to differentiate types of database objects. For example, start indexes with *I* or *X* and databases with *D*. For tables and similar objects, though, as discussed earlier, this approach is inappropriate.

In general, do not impose unnecessary restrictions on the names of objects accessed by end users. Relational databases are supposed to be user friendly. A strict database naming convention, if not developed logically, can be antithetical to a useful and effective database environment. Some

organizations impose arbitrary length limitations on database tables, such as an 8-byte limit even though the DBMS can support up to 128-byte table names. There is no practical reason to impose a limitation on the length of database table names.

Table names should be as descriptive as possible, within reason. Furthermore, the same naming conventions should be used for all “tablelike” objects, including views, synonyms, and aliases, if supported by the DBMS. Each of these objects is basically a collection of data accessible as rows and columns. Developing separate naming conventions for each is of no real value. With this approach, database objects that operate like tables will be defined similarly with a very descriptive name. The type of object can always be determined by querying the DBMS system catalog or data dictionary.

Avoid encoding table names to make them shorter.

Encoding table names to make them shorter is another arbitrary naming standard that should be avoided. Table names should include a 2- or 3-byte application identification prefix, followed by an underscore and then a clear, user-friendly name. For example, a good name for the table containing employee information in a human resources system would be HR_EMPLOYEE. You may want to drop the application identification prefix from the table name for tables used by multiple applications.

Keep in mind, too, that some database object names will, in some cases, be externalized. For instance, most DBMSs externalize constraint names when the constraint is violated. There are many types of constraints—triggers, unique constraints, referential constraints, check constraints—each of which can be named. Keeping the names consistent across environments allows the error messages to be consistent. If the DBMS delivers the same error message in the development, test, integration, and production environments, debugging and error correction will be easier.

Standard Abbreviations

Although you should keep the database object names as English-like as possible, you will inevitably encounter situations that require abbreviations. Use abbreviations only when the full text is too long to be supported as an object name or when it renders the object name unwieldy or difficult to remember. Create a list of standard abbreviations and forbid the use of

nonstandard abbreviations. For example, if “ORG” is the standard abbreviation for “organization,” do not allow variants such as “ORGZ” to be used. Using standard abbreviations will minimize mistyping and make it easier for users to remember database object names. Adhering to this practice will make it easier to understand the database objects within your environment.

Create a list of standard abbreviations.

Other Database Standards and Procedures

Although database naming standards are important, you will need to develop and maintain other types of database standards. Be sure to develop a comprehensive set of standards and procedures for each DBMS used by your organization. Although you can write your database standards from scratch, there are other potentially easier ways to build your standards library. Basic standards that can be modified to your requirements can be bought from a publisher or software vendor. Or you can gather suggested standards from the community via user groups and conferences.

Regardless of whether they are purchased, written in house, or adopted from a user group or committee, each of the following areas should be covered.

Roles and Responsibilities

The successful operation of a DBMS requires the coordinated management efforts of many skilled technicians and business experts. A matrix of database management and administration functions should be developed that documents each support task and who within the organization provides the support. The matrix can be created at a departmental level, a job description level, or even by individual name. A sample matrix is shown in [Table 2.4](#). An *X* in the matrix indicates involvement in the process, whereas a *P* indicates primary responsibility.

Table 2.4. Database Support Roles and Responsibilities

Task	DBA	DA	SA	Management	Operations	Applications	End Users
DBMS budget	X		X	P		X	X
DBMS installation	P		X		X	X	X
DBMS upgrade	P		X	X	X	X	X
Database usage policy	P	X		X			
Capacity planning	X		P	X	X	X	
Data modeling and analysis	X	P					X
Metadata policy	X	P		X			X
Governance and compliance	X	X	X	X			P
Database design	P	X				X	
Database creation	P						
System performance	X		P				
Database performance	P		X			X	
Application performance	X		X			P	
Backup and recovery	P		X		X	X	
Disaster recovery	P		X		X		
Database security	P		X		X		
Stored procedures	X					P	
Triggers	P					X	
User-defined functions	X					P	
Application design	X					P	
Application turnover	X				X	P	X
Application design reviews	X	X	X	X	X	P	X

Of course, you can create whatever tasks you deem necessary in your roles and responsibilities matrix. You may need additional tasks, or fewer than in this sample. For example, you may wish to differentiate between stored-procedure development, testing, and management by creating a different task category for each and breaking down the support requirements differently.

Whatever the final format of your roles and responsibilities matrix, be sure to keep it accurate and up-to-date with new DBMS features and tasks. An up-to-date matrix makes it easier to define roles within the organization and to effectively apportion database-related workload.

Communication Standards

You might also choose to develop specific standards for communication between groups and specific personnel. For example, you might want to document how and when the DBA group must communicate with the systems programming group when a new DBMS release is being installed.

Developing robust communication standards can simplify a DBA's job during the inevitable downtime that occurs due to system, application, or even hardware errors. For example, consider adopting a standard whereby the DBA communicates only with the manager during troubleshooting and

emergency remediation. This keeps the manager informed and enables the DBA to dodge the dozens of phone calls that come in from angry users, the help desk, and so on. The manager can communicate the status outward while the DBA focuses exclusively on troubleshooting and getting the systems back up and running again.

Data Administration Standards

If a DA group exists within your organization, they should develop a basic data administration standards guide to outline the scope of their job role. If a DA group does not exist, be sure to include DA standards in the DBA standards as appropriate.

Include DA standards in the DBA standards as appropriate.

The data administration standards should include the following items:

- A clear statement of the organization's overall policy with regard to data, including its importance to the company
- Guidelines for establishing data ownership and stewardship
- Rules for data creation, data ownership, and data stewardship
- Metadata management policy
- Conceptual and logical data modeling guidelines
- The organization's goals with regard to creating an enterprise data model
- Responsibility for creating and maintaining logical data models
- Guidelines for tool usage and instructions on how data models are to be created, stored, and maintained
- Organizational data-sharing policies
- Instructions on how to document when physical databases deviate from the logical data model
- Guidelines on communication between data administration and database administration to ensure effective database creation and usage

Database Administration Standards

A basic set of database administration standards should be established to ensure the ongoing success of the DBA function. The standards serve as a guide to the DBA services offered and to specific approaches to supporting

the database environment. For example, standards can be developed that outline how requests are made to create a new database or make changes to existing databases, and that specify which types of database objects and DBMS features are favored and under which circumstances they are to be avoided. Standards can establish backup and recovery procedures (including disaster recovery plans) and communicate the methods used to transform a logical data model into a physical database implementation. An additional set of DBA standards that cover database performance monitoring and tuning may be useful to document procedures for overcoming performance problems.

The DBA standards serve as a guide to specific approaches to supporting the database environment.

Although the DBA standards will be most useful for the DBA staff, the application development staff will need them to learn how best to work with the DBA staff. Furthermore, any performance tuning tricks that are documented in the DBA standards should be shared with programmers. The more the application programmers understand the nuances of the DBMS and the role of the DBA, the better the working relationship between DBA and development will be—resulting in a more efficient database environment.

System Administration Standards

Once again, standards for system administration or systems programming are required only if your organization separates the SA function from the DBA function. System administration standards are needed for many of the same reasons that DBA standards are required. Standards for SA may include

- DBMS installation and testing procedures
- Upgrade policies and procedures
- Bug fix and maintenance practices
- A checklist of departments to notify for impending changes
- Interface considerations
- DBMS storage, usage, and monitoring procedures

Database Application Development Standards

The development of database applications differs from typical program development. You should document the special development considerations

required when writing programs that access databases. The database application development standards should function as an adjunct to any standard application development procedures within your organization. This set of standards should include

- A description of how database access differs from flat file access
- SQL coding standards
- SQL performance tips and techniques
- Program preparation procedures and guidance on how to embed SQL in an application program
- Interpretations of SQLSTATEs and error codes
- References to other useful programming materials for teleprocessing monitors, programming languages, and general application development standards

Database Security Standards

The DBA group often applies and administers DBMS security. However, at some shops the corporate data security unit handles DBMS security. A resource outlining the necessary standards and procedures for administering database security should contain the following information:

- Details on what authority to grant for specific types of situations; for example, if a program is being migrated to production status, what DBMS authorization must be granted before the program will operate successfully in production
- Specific documentation of any special procedures or documentation required for governance- and compliance-related requests
- A definitive list of who can approve what types of database authorization requests
- Information on any interfaces being used to connect DBMS security with operating system security products
- Policies on the use of the WITH GRANT OPTION clause of the SQL GRANT statement and how cascading REVOKES are to be handled
- Procedures for notifying the requester that database security has been granted
- Procedures for removing security from retiring, relocating, and

terminated employees

Outline necessary standards and procedures for administering database security.

Application Migration and Turnover Procedures

As discussed earlier, the minimum number of environments for supporting database applications is two: test and production. Some organizations, however, create multiple environments to support, for example, different phases of the development life cycle, including

- *Unit testing*—for developing and testing individual programs
- *Integration testing*—for testing how individual programs interoperate
- *User acceptance testing*—for end user testing prior to production status
- *Quality assurance*—for shaking out program bugs
- *Education*—for training end users how to work the application system

When multiple environments exist, procedures are required for migrating database objects and programs from environment to environment. Specific guidelines are needed to accomplish migration in a manner conducive to the usage of each environment. For example, what data volume is required for each environment and how is data integrity to be assured when testing activity occurs? Should data be migrated, or just the database structures? How should existing data in the target environment be treated—should it be kept, or overlaid with new data? Comprehensive migration procedures should be developed to address these types of questions.

Procedures are required for migrating database objects and programs from environment to environment.

The migration and turnover procedures should document the information required before any database object or program can be migrated from one environment to the next. At a minimum, information will be required about the requester, when and why the objects should be migrated, and the appropriate authorization to approve the migration. To ensure the success of the migration, the DBA should document the methods used for the migration and record the verification process.

Design Review Guidelines

All database applications should be subjected to a design review at various stages of their development. Design reviews are important to ensure proper application design, construction, and performance. Design reviews can take many forms. [Chapter 6, “Design Reviews,”](#) offers a comprehensive discussion.

Operational Support Standards

Operational support is defined as the part of the IT organization that oversees the database environment and assures that applications are run according to schedule. Sufficient operational support must be available to administer a database environment effectively. The operational support staff is usually the first line of defense against system problems. Program failures, hardware failures, and other problems are first identified by operational support before specialists are called to resolve the problems.

Operational support assures that applications are run according to schedule.

Standards should be developed to ensure that the operational support staff understands the special requirements of database applications. Whenever possible, operational support personnel should be trained to resolve simple database-related problems without involving the DBA because the DBA is a more expensive resource.

DBMS Education

Organizations using DBMS technology must commit to ongoing technical education classes for DBAs, programmers, and system administrators. Provide a catalog of available courses covering all aspects of DBMS usage. At a minimum, the following courses should be made available:

- *DBMS Overview*: a one-day management-level class that covers the basics of DBMS
- *Data Modeling and Database Design*: a thorough course covering conceptual, logical, and physical database design techniques for DAs and DBAs
- *Database Administration*: in-depth technical classes for DBAs, SAs, and systems programmers
- *Introduction to SQL*: an introductory course on the basics of SQL for

every DBMS user

- *Advanced SQL*: an in-depth course on complex SQL development for DBAs and programmers
- *Database Programming*: an in-depth course for application programmers and systems analysts that teaches students how to write programs that use the DBMS

Each of these courses should be available for each DBMS installed in your organization. Furthermore, provide training for any other database-related functionality and software such as proper use of database utilities, query and reporting tools, and DBA tools.

Commit to ongoing technical education classes.

DBMS education can be delivered using a variety of methods, including instructor-led courses, computer-based training, Web-based training, and distance learning. Sources for DBMS education include DBMS vendors, ISVs, consultants (large and small, international and local), and training specialists (such as Themis and ProTech).

Finally, be sure to make the DBMS reference material available to every user. Most vendors offer their DBMS reference manuals in an online format using Adobe Acrobat files or Windows Help. Be sure that each user is given a copy of the manuals or that they are available in a central location to minimize the amount of time DBAs will have to spend answering simple questions that can be found in the DBMS documentation.

Summary

Comprehensive advance planning is required to create an effective database environment. Care must be taken to select the correct DBMS technology, implement an appropriate DBMS upgrade strategy, develop useful database standards, and ensure the ongoing availability of education for database users. By following the guidelines in this chapter, you can achieve an effective database environment for your organization.

Nevertheless, setting up the database environment is only the beginning. Once it is set up, you will need to actively manage the database environment to ensure that databases are created properly, used correctly, and managed for performance and availability. Read on to discover how the DBA can

accomplish these tasks.

Review

1. Why should database standards be implemented and what are the risks associated with their lack?
2. What are the potential risks of upgrading to a new DBMS release without a plan?
3. What is the difference between a version and a release of a DBMS?
4. Name the three TPC benchmarks and describe how they differ from one another.
5. Describe the four levels of DBMS architecture in terms of the type and nature of processing to which each is best suited.
6. What are the factors to be considered when calculating total cost of ownership (TCO) for a DBMS?
7. Name five requirements that must be planned for when installing a new DBMS.
8. Describe the difference between software that supports a DBMS release and software that exploits a DBMS release.
9. How many standard abbreviations should be supported for a single term? Why?
10. What is wrong with the following SQL code for creating a relational table? (Do not approach this question from a syntax perspective; consider it, instead, in terms of database naming standards.)

```
CREATE TABLE tg7r5u99_p
(c1    INTEGER NOT NULL,
 c2    CHAR(5)  NOT NULL,
 c9    DATE)
;
```

Bonus Question

Your DBMS vendor, MegaDataCorp, just announced the general availability of the latest and greatest version of MDC, the DBMS you use. MDC Version 9 supports several new features that your users and developers have been clamoring for over the past year. You are currently running MDC Version 7.3. Prepare a short paper discussing your plans for upgrading to MDC

Version 9, and outline the potential benefits and risks of your upgrade plan.

Suggested Reading

Blaha, Michael R. *A Manager's Guide to Database Technology*. Upper Saddle River, NJ: Prentice Hall (2001). ISBN 0-13-030418-2.

Connolly, Thomas, and Carolyn Begg. *Database Systems: A Practical Approach to Design, Implementation, and Management*. 4th ed. Harlow, England: Addison-Wesley (2004). ISBN 978-0-321-29401-2.

3. Data Modeling and Normalization

Data modeling is the process of analyzing the things of interest to your organization and how these things are related to each other. The data modeling process results in the discovery and documentation of the data resources of your business. Data modeling asks the question “What?” instead of the more common data-processing question “How?”

Data modeling asks the question “What?”

Before implementing databases of any sort, a DBA needs to develop a sound model of the data to be stored. Novice database developers frequently begin with the quick-and-dirty approach to database implementation. They approach database design from a programming perspective. Because novices often lack experience with databases and data requirements gathering, they attempt to design databases like the flat files they are accustomed to using. This is a major mistake. Indeed, most developers using this approach quickly discover problems after the databases and applications become operational in a production environment. At a minimum, performance will suffer and data may not be as readily available as required. At worst, data integrity problems and/or performance problems¹ may arise, rendering the entire application unusable.

A proper database design cannot be thrown together quickly by novices. What is required is a practiced and formal approach to gathering data requirements and modeling the data, that is, the discovery and identification of entities and data elements. Data normalization is a big part of data modeling and database design. A normalized data model reduces data redundancy and inconsistencies by ensuring that the data elements are designed appropriately.

It is actually quite simple to learn the basics of data modeling, but it can take a lifetime to master all its nuances. This chapter introduces the concepts of data modeling and normalization and provides some general guidelines to their proper use.

Experienced data modelers and DAs will probably find the material in this chapter incomplete. This is by design: The intent is to introduce the concept

to DBAs and guide them down the right path to producing databases from properly developed data models. A complete treatment of data modeling requires an entire text. Numerous books exist that provide an exhaustive treatment of data modeling and normalization. Consult the suggested references at the end of this chapter for additional details on data modeling.

Data Modeling Concepts

A popular folk story about four blind men and an elephant helps to illustrate the purpose of data modeling:

Four blind men happened upon an elephant during the course of their journey. The blind men had never encountered an elephant before, but they were a curious group. Therefore, each blind man attempted to learn about the elephant by touching it. The first blind man grabbed the elephant by the trunk and exclaimed, “Oh! An elephant is like a snake—a long, slithery tube.” The second blind man reached out, touched the side of the elephant, and remarked, “No, no, the elephant is more like a wall—very flat and solid.” The third blind man was confused, so he reached out to touch the elephant but poked his hand on a tusk and said, “No, you’re both wrong, the elephant is more like a spear than anything else!” The fourth blind man grabbed the elephant by the leg and shouted, “You’re all wrong, an elephant is very much like a tree—round and solid.”

Well, each blind man was right, but he was also wrong. The problem was not with the experience of each blind man, but with the scope of that experience. To be a successful data modeler you must learn to discover the entire “truth” of the data needs of your business. You cannot rely on the requirements of a single user, or even a single expert, because that person’s scope of experience will not be comprehensive. The goal of a data model is to record the data requirements of a business process. The scope of the data model for each line of business must be comprehensive. If an enterprise data model exists for the organization, each individual line-of-business data model must be verified for correctness against the overall enterprise data model.

The goal of a data model is to record the data requirements of a business process.

Data modeling, at the most basic level, begins as a conceptual venture. The first objective of conceptual data modeling is to understand the requirements. A data model, in and of itself, is of limited value. Of course, a data model delivers value by enhancing communication and understanding, and it can be argued that these are quite valuable. However, the primary value of a data model is its ability to be used as a blueprint to build a physical database.

When databases are built from a well-designed data model, the resultant structures provide increased value to the organization. The value derived from the data model exhibits itself in the form of minimized data redundancy, maximized data integrity, increased stability, better data sharing, increased consistency, more timely access to data, and better data usability. These qualities are achieved because the data model clearly outlines the data resource requirements and relationships in a clear, concise manner. Building databases from a data model will result in a better database implementation because you will have a better understanding of the data to be stored in your databases.

The Enterprise Data Model

An [*enterprise data model*](#) is a single data model that describes comprehensively the data needs of the entire organization. This chapter discusses data models as they pertain to individual projects or applications, but not the relatively advanced topic of the enterprise data model.

Managing and maintaining an enterprise data model is fraught with many non-database-related distractions such as corporate politics and ROI that are hard to quantify.

If your organization has committed to developing and maintaining an enterprise data model, you should definitely take the time to read more about the topic by consulting the books referenced at the end of this chapter.

Another benefit of data modeling is the opportunity to discover new uses for data. A data model can clarify data patterns and potential uses for data that might otherwise remain hidden. Discovery of such patterns can change the way your business operates and can potentially lead to a competitive advantage and increased revenue for your organization.

A data model can clarify data patterns and potential uses for data.

Data modeling requires a different mindset from requirements gathering for application development and process-oriented tasks. It is important to think “what” is of interest instead of “how” tasks are accomplished. To transition to this alternative way of thinking, follow these three rules:

- *Don't think physical; think conceptual.* Concern yourself with business issues and terms rather than physical storage issues and DBMS constraints.
- *Don't think process; think structure.* How something is done, although important for application development, is not important for data modeling. The things to which processes are being applied are what are important to data modeling.
- *Don't think navigation; think relationship.* The way that things are related to one another is important, because relationships map the data model. The way in which relationships are traversed is unimportant to conceptual and logical data modeling.

Keep in mind that as you create your data models, you are developing the lexicon of your organization's business. Much as a dictionary functions as the lexicon of a given language, the data model functions as the lexicon of business terms and their usage.

Entity-Relationship Diagramming

Data models are typically rendered in a graphical format using an *entity-relationship diagram*, or E/R diagram for short. An E/R diagram graphically depicts the entities and relationships of a data model. [Figure 3.1](#) shows a sample E/R diagram. This diagram was produced using ERwin, from Computer Associates, a leading data modeling and design tool. Many popular data modeling tools are available from a variety of vendors. Consult [Appendix C](#) for a list of data modeling tools and vendors.

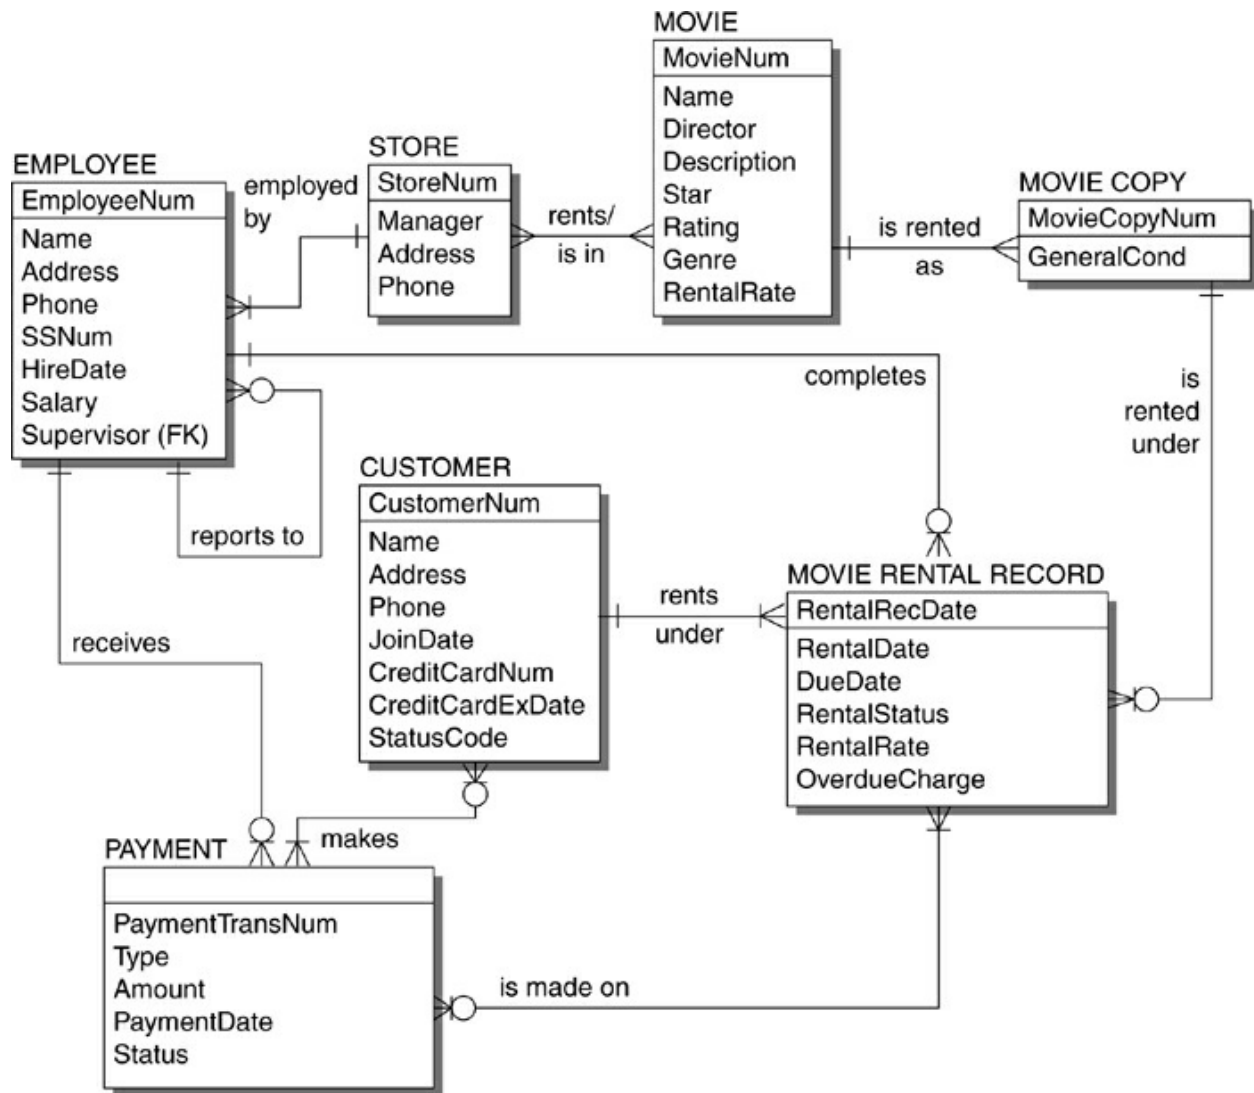


Figure 3.1. A sample E/R diagram

An E/R diagram graphically depicts the entities and relationships of a data model.

Multiple diagramming methods and techniques are available for E/R diagramming. Regardless of the method used, an E/R diagram will show the entities in boxes and use lines or arrows to show relationships. The format of the boxes and lines differs with the diagramming method. Several of the most common E/R diagramming methods are shown in [Figure 3.2](#). There is nothing that really makes one method superior to the others, so choose one and try to stick to it to minimize confusion. Allowing different modeling teams to use different diagramming techniques will make sharing the

knowledge contained in the data models more difficult. Furthermore, it complicates the consolidation of the data models into an enterprise data model, if one is maintained.

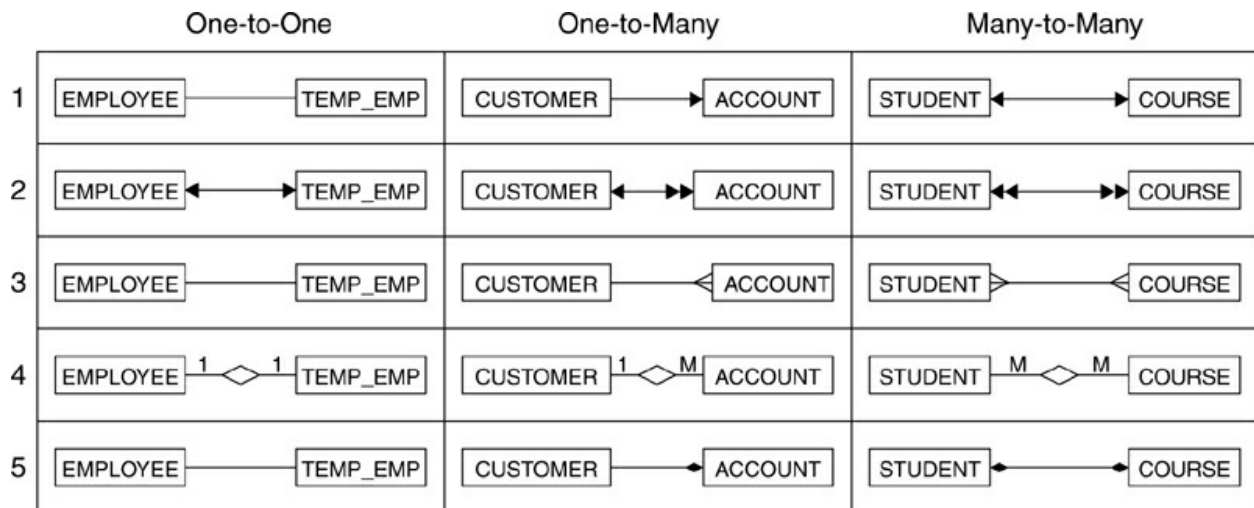


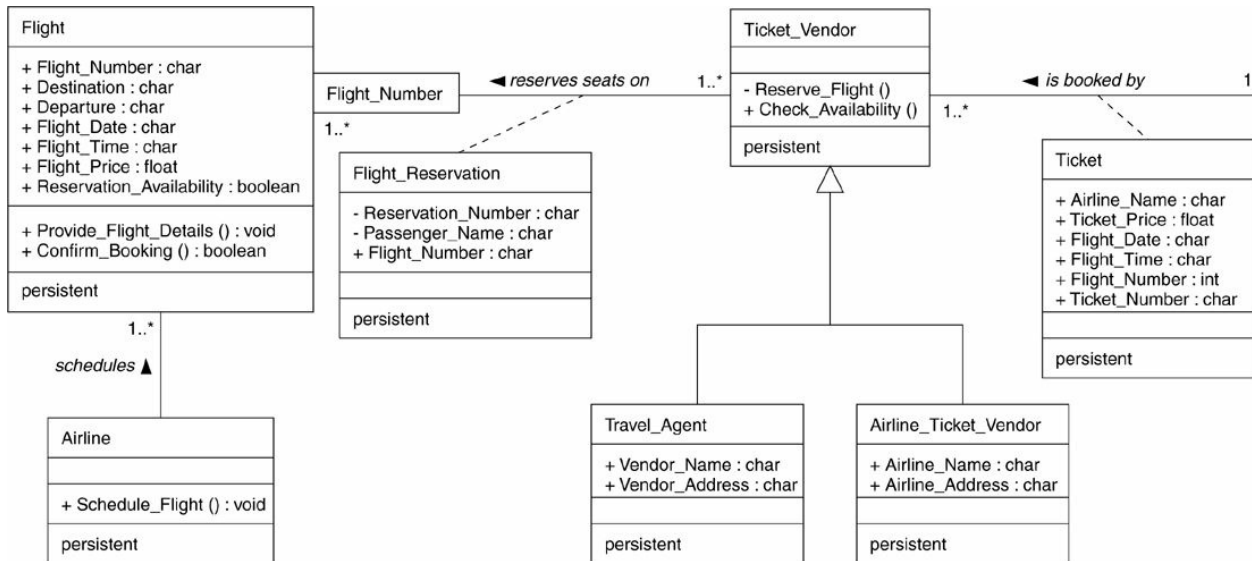
Figure 3.2. E/R diagramming methods: (1) Ross, (2) Bachmann, (3) Martin, (4) Chen, (5) Rumbaugh

Refer to [Figure 3.2](#) again. The major difference among the diagramming methods is the way in which cardinality is shown. *Cardinality* refers to the maximum number of instances an entity can take on. Each row of [Figure 3.2](#) shows how a different E/R diagramming style handles one-to-one, one-to-many, and many-to-many relationships. The E/R diagramming methods are named after the originator of the technique. By row, the names of the diagramming techniques depicted in [Figure 3.2](#) are Ross, Bachmann, Martin, Chen, and Rumbaugh.

Most modeling tools support one or more of these diagramming techniques. The most popular diagramming techniques are Martin (also known as Information Engineering) and Ross. The E/R diagram in [Figure 3.1](#) was developed using the Martin technique. Popularity aside, each of these techniques is used by many data modelers, and you should select the technique that is easiest for those in your organization to use and understand.

Another popular modeling technique is Unified Modeling Language. UML is a consolidation of several popular object-oriented notations and concepts. It originated from the work of Grady Booch, James Rumbaugh, and Ivar Jacobson. UML provides a standard set of diagrams and notations for modeling object-oriented systems. UML actually defines nine types of

diagrams for modeling different types of systems. The UML class diagram can be used to model data, but traditional E/R diagrams capture more information about pure data and are better suited than UML to data modeling and relational database design tasks. [Figure 3.3](#) shows a sample UML class diagram depicting a data model for an airline. Because objects comprise both data and process, UML is more than a data modeling diagramming technique.



Courtesy Lou Varveris, IBM. Diagram created with IBM Rational System Architect.

Figure 3.3. Sample UML class diagram

The Components of a Data Model

A *data model* is built using many different components that act as abstractions of real-world things. The simplest data model consists of entities and relationships. As work on the data model progresses, additional detail and complexity are added. Let's examine the many different components of a data model and the terminology used for data modeling.

Entities

An *entity*, at a very basic level, is something that exists and is capable of being described. It is a person, place, thing, concept, or event about which your organization maintains facts. For example, STUDENT, INSTRUCTOR, and COURSE are specific entities about which a college or university must be knowledgeable to perform its business.

An entity is a person, place, thing, concept, or event.

Entity Naming Guidelines

It is important to follow a rigorous standard for naming entities within a data model. If developers attempt to communicate using poorly defined and nonstandard names, confusion will result. A good standard to follow is to always name entities in singular form using capital letters. An entity should be viewed as a pattern for the occurrences therein, instead of the set of all occurrences for that type of entity. For example, an entity that contains data about your company's employees should be named EMPLOYEE rather than EMPLOYEES.

Furthermore, an entity should always be a noun, or a combination of an adjective with a noun. However, use of adjectives should be minimized whenever possible. At times, the adjective can be an attribute in disguise. For example, an entity named CONTRACT EMPLOYEE is better expressed as an entity named EMPLOYEE with an attribute describing the status of the employee (where the value of Status can be FULLTIME or CONTRACT).

An entity should always be a noun, or a combination of an adjective with a noun.

Promote consistency in the choice of terminology used for entities. In general, use the term favored by the business users. When the business uses multiple terms (synonyms) for the same entity and there is no clear consensus, choose one and use it consistently. For example, if some users prefer VENDOR but others prefer SUPPLIER, choose one and use it consistently throughout the data model. In some cases, common sense may dictate a choice when the business cannot express a preference. For example, CLIENT is perhaps a better term than USER, given the negative connotation sometimes given this word ("drug user").

Finally, remove process-specific artifacts from the entity name. For example, AGENT is better than SELLING AGENT. Remember, the data model should capture the "what," not the "how," and "selling" describes a process.

Entity Occurrences

An [*occurrence*](#) is an instance of an entity. For example, AUTHOR is an example of an entity. Craig S. Mullins, along with all my descriptive details,

is an occurrence of the AUTHOR entity. It is important to differentiate the abstract entity from the concrete occurrences of the entity. *Entity instance* is an equivalent term for occurrence.

An occurrence is an instance of an entity.

Attributes

An *attribute* is a characteristic of an entity. Every attribute does one of three things:

Every attribute does one of three things: It identifies, relates, or describes.

- *Identifies*. An attribute that identifies is a *candidate key*. If the value of an identifying attribute changes, it should identify a different entity occurrence. An attribute that identifies should be immutable.
- *Relates*. An attribute that relates entities is a *foreign key*. The attribute refers to the primary key attribute of an occurrence of another (or the same) entity.
- *Describes*. An attribute is descriptive if it depicts or expresses a characteristic of an entity occurrence but does not identify or relate.

As a rule of thumb, nouns tend to be entities and adjectives tend to be attributes. Nevertheless, this is not a hard-and-fast rule: Be sure to apply knowledge of the business to determine which nouns and attributes are entities and which are attributes. Every attribute must identify the entity occurrence, relate the entity occurrence to another entity occurrence (in the same or another entity), or describe the entity occurrence.

An attribute must definitively reflect its specific, intended meaning. The instances of an attribute must be atomic in nature—that is, an attribute represents a singular fact that cannot be further decomposed. For that reason, using Name as an attribute for, say, the entity PERSON would not be good practice because Name can be decomposed into FirstName, Middle-Initial, and LastName.

Domains

Each attribute is assigned a valid domain. The *domain* defines the universe of valid values for the data element. An example of a domain would be “valid

positive integer values between 1 and 10.” *Data type* is a component of the domain. The data type specifies, appropriately enough, the type of data that comprises the domain; some examples are Integer, Decimal, Character, Date, Time, and so on.

The domain defines the universe of valid values for the data element.

When an attribute is defined that represents a code, the domain for the code should, if possible, consist of self-documenting values. For example, the domain for an attribute defining a periodic frequency of weekly, monthly, quarterly, and yearly is better defined as (“W,” “M,” “Q,” “Y”) than as (“1,” “2,” “3,” “4”). The first domain is self-documenting and easy to remember. The second is not so easy to remember.

Attribute Naming Guidelines

Develop and adhere to a standard format for naming attributes. For example, if you name the attributes in singular form with InterCaps (“FirstName”), don’t switch back and forth to all lowercase (“first_name”). Choose one form and stick to it. Some potential attribute names for the EMPLOYEE entity might be StreetAddress (or street_address), ZipCode (or zip_code), and EmployeeId (or employee_id). The point is not to spend a lot of time agonizing over the exact standard for naming attributes, but to choose a standard and stick to it to minimize confusion and clarify communication.

Adhere to a standard format for naming attributes.

An *attribute name* should consist of a prime descriptive word coupled with a class word. A class describes a particular subset of attributes that are common to most information systems. Additional qualifying words can be used if necessary to define the attribute accordingly. The prime descriptive word can be the entity name or another descriptive word. For the class word, establish a valid list of classes for the types of domains in use. A sample list of possible classes is shown in [Table 3.1](#). You may decide to use the entire class name or a standard abbreviation. The following are some sample attribute names:

- *VendorID*—where Vendor is the entity, and the attribute is an identifier.

- *ProductName*—where Product is the entity, and the domain of the attribute is an alphabetic, character name.
- *ProductYearlySalesAmount*—where Product is the entity, and the domain is an amount. However, there may be multiple amounts for each product instance. This particular attribute is further qualified by Yearly, which indicates the period, and Sales, which indicates that the amount is an amount sold, as opposed to, say, an amount purchased.

Table 3.1. Sample Attribute Class Words

Class	Class Description	Class Abbreviation
Address	Street address or location	ADDR
Amount	Monetary amount	AMT
Audio	Digital audio or voice recordings	AUD
Code	Codes, types, and classifications	CDE
Date	Calendar dates	DAT
Description	Descriptive text	DSC
Identifier	Unique identifier	ID
Image	Digital images	IMG
Name	Alphabetic, character name	NME
Number	Numeric counter	NUM
Percentage	A numeric ratio or percentage	PCT
Quantity	Numeric units or amounts (nonmonetary)	QTY
Text	Free-format text document	TXT
Time	Clock time	TME
Video	Digital video streams; moving pictures	VID

Homonyms and Synonyms

According to *Merriam-Webster's Dictionary of English Usage* (online):

Homonym, noun

1: one of two or more words spelled and pronounced alike but different in meaning (as the noun *quail* and the verb *quail*)

Synonym, noun

1. one of two or more words or expressions of the same language that have the same meaning or nearly the same meaning in some or all senses

Using homonyms and synonyms in your data modeling can cause confusion. A homonym must be used in context for its meaning to be understood. When two words are spelled and pronounced identically, the only way to differentiate them is to examine how each is being used. For example, if I just say the word *watch*, am I referring to a timepiece, or am I requesting that you look at something? Devoid of context, it is impossible to tell.

Synonyms cause confusion for a different reason. When the same thing is referred to in different ways, it can become unclear that more than one word is being used to refer to the same thing. For example, are clients and customers the same thing?

Although homonyms and synonyms cannot be banned during the course of conducting business, they can be banned within the scope of data modeling and database design. Be sure that as a data modeler you do so. However, don't attempt to force the line-of-business users to adopt the terminology in place of their current usage habits. Instead, be consistent within the data model and allow the users to continue their current practices without demanding a change of their lexicon. Nevertheless, be sure to capture and document all business usage of homonyms and synonyms.

Your list of class words may differ. On the other hand, the list might be the same, with different abbreviations. Either is fine, as long as you develop a list of class words and enforce their usage. Class words should not have multiple meanings. One, and only one, class word should be defined for each domain. Also, avoid homonyms and synonyms when defining your class words (see the sidebar). Keep in mind that these class words are important for the logical data model only and might not be used in the physical database due to sizing limitations for DBMS object names.

Class words should not have multiple meanings.

Attribute Values

An *attribute value* is an instance of an attribute. The value "Craig" is an

example of an attribute value for the FirstName attribute within the AUTHOR entity. It is important to differentiate the abstract attribute from the concrete values for the attribute.

Nulls and Lack of Value

A [*null*](#) represents missing or unknown information at the attribute level. If an attribute “value” can be null, it means one of two things: The attribute is not applicable for certain occurrences of the entity, or the attribute applies to all entity occurrences, but the information may not always be known. Of course, it could be a combination of these two situations, too.

A null represents missing or unknown information.

For example, suppose that an entity, STUDENT, contains an attribute HairColor. Further, the HairColor attribute can be null. But what does a null HairColor indicate? Consider three potential entity occurrences: a man with black hair, a woman with unknown hair color, and a bald man. Both the woman with the unknown hair color and the bald man could be assigned a null HairColor, but for different reasons. The woman’s HairColor would be null, meaning “presently unknown”; the bald man’s HairColor would be null, meaning “not applicable.”

Of course, the HairColor of a bald man can be expressed by extending the domain of colors to include the value “bald.” But what of a domain that is not feasibly extensible? For example, consider a Date such as a TerminationDate attribute in an EMPLOYEE entity. The domain of a Date is potentially all valid calendar dates. It might not be practical to assign a special date value to mean “unknown.” Consider the value you would like in the TerminationDate of your EMPLOYEE occurrence. Null, meaning unknown, is far preferable to a valid date, regardless of how far in the future, isn’t it?

“Not Applicable” Nulls May Signal a Design Problem

Using null to indicate “not applicable” can be an indication of improper data design. By properly modeling and normalizing your data structures, you can usually eliminate the need to use nulls to indicate that a column is inapplicable for a specific row.

A properly normalized data model should result in each attribute value depending upon the primary key of the entity occurrence. If a

particular attribute is “not applicable,” arguably the design is improperly normalized. For more information on normalization consult the information in this chapter beginning with the section “[What Is Normalization?](#)”

Be sure to consider the nuances of nulls as you develop your data model. Be sure to determine the nullability for every attribute.

Nulls are sometimes inappropriately referred to as null values. Using the term *value* to describe a null is inaccurate because the term *null* implies the *lack* of a value.

Keys

A *key* consists of one or more attributes, the values of which uniquely identify an entity occurrence and define relationships between entities. Well, more precisely, candidate keys and primary keys identify the entity. A combination of the primary key value of one entity and the foreign key value of another entity identifies a relationship. For example, if CustNo is the primary key of the CUSTOMER entity, CustNo will be used as the foreign key in a related entity, such as CUSTOMER_CONTACT. A customer can have multiple contacts, so the same CustNo would be registered for each contact for the customer.

Keys consist of the attributes that identify entity occurrences and define the relationships between entities.

A key should contain no embedded meaning. The key’s purpose is to identify, not to describe. The other attributes in the entity serve a descriptive purpose. When a key contains an embedded meaning, problems can arise if the meaning changes. Furthermore, the values for any embedded meaning are likely to be outside your immediate control, which can also cause data integrity and modification problems.

Candidate Keys

Each entity can have multiple candidate keys, but it must have at least one. A *candidate key* is an attribute, or set of attributes, that can be used to uniquely identify an occurrence of the entity. If the value of the attribute cannot be used to identify a specific occurrence of the entity, it does not represent a

candidate key.

Primary Key

Each entity will have one, and only one, primary key. The [*primary key*](#) is chosen from the set of candidate keys and is used to identify an entity occurrence. Choosing an appropriate primary key is important because the primary key will be used to define the foreign keys in related, dependent entities.

Characteristics of good primary keys include the following:

- The primary key must guarantee the uniqueness of an entity occurrence.
- The value of any component of the primary key cannot be null.
- Primary keys of basic entities should not have embedded meaning.
- Primary keys should be immutable—that is, not capable of or susceptible to change.

Furthermore, you should have control over primary key values. When values are assigned externally, you lose control, causing potential data problems. For one, it is impossible to ensure that the key values will always be unique. As an example, a Social Security number is a bad primary key choice to identify employees because it is assigned outside of your control. A numeric identifier assigned by your organization is a better choice.

Foreign Keys

[*Foreign keys*](#) reside in dependent entities to establish relationships. The primary key identifies an entity occurrence; foreign keys identify relationships between entity occurrences. For one-to-many relationships, the foreign key will always be on the many side of a relationship. For one-to-one relationships, determination of foreign key placement is more difficult. The basic rule is to analyze the attributes and assign the key to the entity for which it is most characteristic. If one end of the one-to-one relationship is optional, the entity at the optional end should contain the foreign key and the entity at the mandatory end should contain the primary key.

Relationships

[*Relationships*](#) define how the different entities are associated with each other. A relationship name should describe the role played by an entity in its

association with another (or perhaps the same) entity. The keys define a relationship: the primary key in the parent entity and the foreign key in the dependent entity.

Relationships define how the different entities are associated with each other.

Relationships are not just the “lines” that connect entities but provide meaning to the data model and should be assigned useful names. The relationship name makes a factual statement about the association between entities. For example, consider the one-to-many relationship between the COURSE and INSTRUCTOR entities: Each COURSE is taught by one, and only one, INSTRUCTOR, but an INSTRUCTOR teaches many COURSEs. The name of the relationship, in this case “is-taught-by,” coupled with the names of the entities participating in the relationship should form a meaningful sentence.

In an E/R diagram, the relationships should be read clockwise from right to left over the line. Following this convention will assure the legibility of your data models.

Cardinality

Cardinality is the number of occurrences that can exist between a pair of entities. Another way of looking at cardinality is as the number of entity occurrences applicable to a specific relationship. Sometimes the term *degree* is used instead of cardinality. Each end of a relationship has a cardinality, or degree, associated with it.

Cardinality is the number of occurrences that can exist between a pair of entities.

At the simplest level, cardinality is expressed by the way relationships are drawn in the E/R diagram. Recall the diagramming techniques displayed in [Figure 3.2](#). The notion of cardinality is expressed as either “one” or “many.” Using the Martin technique (line 3 in the figure), a cardinality of one is expressed as a straight line and a cardinality of many is expressed using crow’s feet.

At a more complex level, the data model should capture more information about the specific cardinality of each relationship. A complete data model

will record the minimum and maximum cardinality for each end of the relationship as integer values. However, such detailed cardinality information need not be represented on the diagram, especially at the conceptual level.

Optionality

The data model must also capture whether relationships are mandatory or optional. This is commonly referred to as the *optionality* of the relationship. Once again, each end of the relationship will have an optionality characteristic.

Each end of the relationship will have an optionality characteristic.

For the Martin diagramming technique, drawing a bar at the end of the relationship indicates that the relationship is mandatory; a small circle indicates an optional relationship. Refer to the relationship diagrammed in [Figure 3.4](#). This data model fragment clearly states that an EMPLOYEE is employed by a STORE. The STORE can have zero, one, or multiple EMPLOYEEs. If an EMPLOYEE exists, a relationship to a STORE is mandatory. Furthermore, an EMPLOYEE can work for only one store.

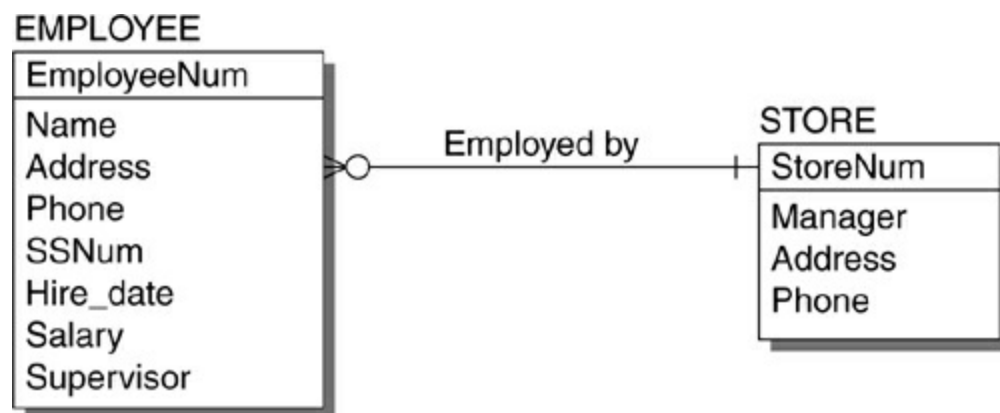


Figure 3.4. Data model fragment

Other diagramming techniques use different methods—perhaps different shapes or specific integer values.

So, you can see how a graphical data model that uses well-defined diagramming techniques and notations can clearly communicate the business requirements of data. Knowing nothing more than data modeling techniques, an individual who is unfamiliar with the business can quickly learn and

communicate the organization's fundamental business rules.

Discovering Entities, Attributes, and Relationships

At a very high level, the data modeling process can be thought of as *the discovery of entities and their attributes along with the relationships between the entities*. The process of discovering and documenting these items for your business is data modeling. Now, data modeling may sound simple enough, but it takes practice to learn how to listen to business speech and be able to identify and differentiate among entities, attributes, and relationships.

Track the nouns, noun phrases, verbs, and adjectives used by the business experts.

The first trick is to keep track of the nouns, noun phrases, verbs, and adjectives used by the business experts. Nouns are potential entities and attributes, or perhaps entity occurrences and attribute values. Adjectives and prepositional phrases are usually potential attributes. Verbs indicate potential relationships.

Of course, developing a data model is more difficult than merely listening for nouns, adjectives, and the like and transforming them into the entities, attributes, and relationships of a data model. Every noun is not necessarily an entity of interest—likewise for adjectives, verbs, and prepositional phrases. You will need to use your experience to filter the business speech to identify the appropriate entities, attributes, and relationships.

When analyzing business descriptions, you will need to place the words in context and ask questions to clearly differentiate between entities and attributes. Remember, attributes identify, relate, or describe entities. Entities are the most important and general “things” the business expert will discuss—the people, places, things, concepts, and events of interest to the business.

Often the terminology used by the business experts will not be appropriate for the data model. You may need to generalize terms, or make them more abstract. For example, business users may speak of doing business with corporations, but business is also conducted with individuals. For the purpose of the data model you may need to create an abstract entity such as PARTY, or some other abstract term that encompasses corporations, other types of organizations, and individuals.

No one individual understands all the data needs of the organization.

When developing a data model, be sure to solicit input from multiple sources. A data model built on the advice of a single expert is sure to be inaccurate. Remember the blind men and the elephant: No one individual understands all the data needs of the organization (or indeed, of any single aspect of the organization). When you talk to multiple experts on multiple occasions, the true nature and definition of the data will become clearer.

Keep in mind that patterns in data models tend to repeat. As you work on data modeling projects, you will begin to notice certain patterns that reappear in multiple subject areas. Learning to recognize these patterns can simplify the creation of data models. However, don't make the mistake of jumping to a pattern that you think you recognize, without consulting multiple business experts. Sometimes what appears to be a recognizable pattern is, in actuality, a completely new pattern.

As with most things, practice will improve your data modeling skills. Don't be discouraged if your initial attempts at creating data models are difficult or error prone. The more you model, the better you will become at data modeling.

Finally, to further clarify the flexibility and usefulness of data modeling, let's look at a data model that incorporates the data modeling terms we just learned. Consult the E/R diagram in [Figure 3.5](#). This E/R diagram shows clearly the relationships among the components of a data model.

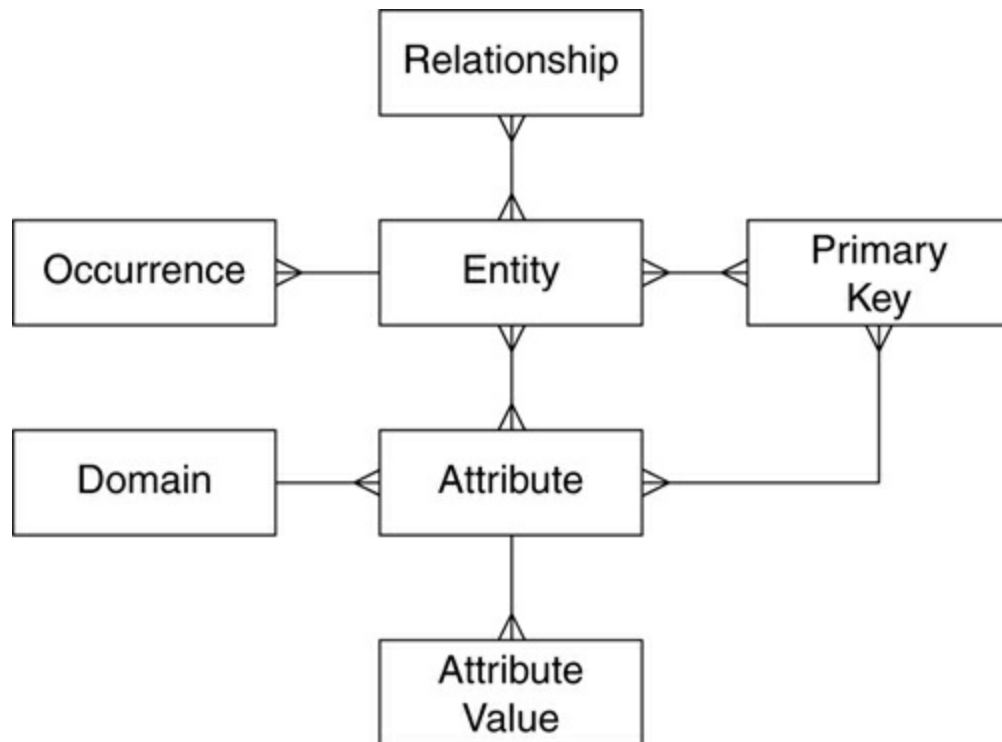


Figure 3.5. Relationships of data modeling terms

Conceptual, Logical, and Physical Data Models

Up until now, we have discussed data modeling in general. But there are really three types of data modeling: conceptual, logical, and physical.

Conceptual and *logical* are loose terms that describe different phases of data modeling. A conceptual data model is generally more abstract and less detailed than a complete logical data model. There is no hard-and-fast boundary separating a conceptual data model from a logical data model. But the following generalities are largely accepted in the modeling community.

A [*conceptual data model*](#) depicts a high-level, business-oriented view of information. Noncritical details can be left out of the conceptual data model in order to emphasize the most important entities, attributes, and relationships. The conceptual data model may contain many-to-many relationships. The goals of conceptual modeling are clarity and simplicity. It is not necessary to discover and document every attribute of each entity at the conceptual level. Furthermore, issues such as cardinality, optionality, and data types can be skipped at the conceptual level. Some simple candidate keys may be recorded, but the conceptual data model will not identify most

keys because the high level of abstraction makes useful key identification impractical or impossible. Furthermore, every relationship need not be named on the conceptual data model. Remember, the goal of a conceptual data model is clarity at a high level. If the relationship names clutter up the data model diagram, they work against that goal.

A conceptual data model depicts a high-level, business-oriented view of information.

A [logical data model](#) consists of fully normalized entities with all attributes defined. Furthermore, the domain or data type of each attribute must be defined. A logical data model requires the specification of candidate keys for unique identification of each occurrence in every entity. For those entities having multiple candidate keys, the logical data model must indicate which candidate key to use for identification, that is, the primary key. Foreign key definitions should be clearly defined or implied by the data relationships. Many-to-many relationships should be translated into associative entities (refer to the sidebar), which may acquire additional attributes and identifiers. Additional details, such as cardinality and whether relationships are optional or mandatory, must be documented in the logical data model. All relationships should be clearly and explicitly named. A logical data model should be a complete document from which a physical database can be developed.

The logical data model offers a comprehensive formal structure that serves as a blueprint for business data.

The [physical data model](#) is created to transform the logical data model into a physical implementation using a specific DBMS product such as DB2, Oracle, or SQL Server. The specifics of the physical data model are covered in [Chapter 4, “Database Design.”](#)

The physical data model transforms the logical data model into a physical implementation using a specific DBMS.

Associative Entities

An [associative entity](#) provides additional information about a relationship, if needed. In other words, an associative entity enables

attributes to be associated with a relationship. Associative entities also make it possible for many-to-many relationships to be implemented using a relational database.

Every many-to-many relationship can be resolved to a pair of one-to-many relationships between each of the two existing entities, and a new entity—an associative entity, as shown in [Figure 3.6](#).

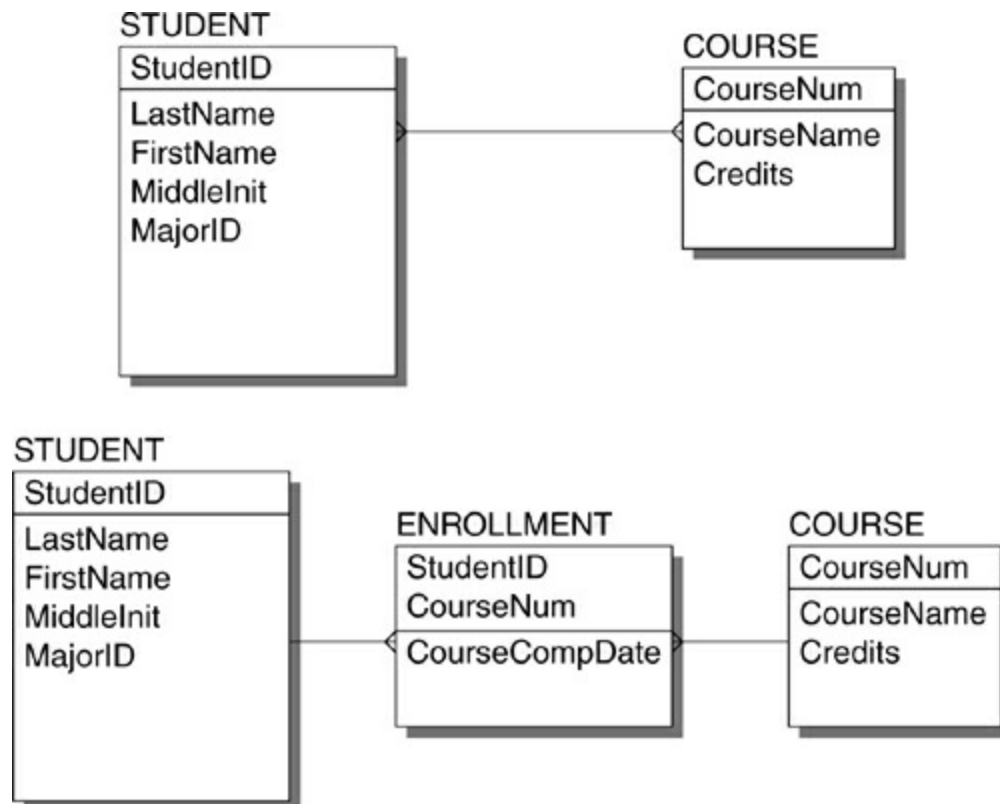


Figure 3.6. From many-to-many to associative entity

In the example shown at the top of the figure, a STUDENT can register for many COURSEs, and a COURSE can contain many STUDENTs. This classic many-to-many relationship can be resolved as shown at the bottom of [Figure 3.6](#). A new associative entity is created, in this case, the ENROLLMENT entity. The primary key from each of the previously existing entities is copied to the new entity. And the many-to-many relationship is replaced with two one-to-many relationships, with the many side, in each case, placed on the new associative entity. Now attributes can be assigned to the new entity, which was not possible for a relationship.

So, a conceptual data model focuses on providing a real-world view of a particular business subject. The conceptual data model makes it easier to understand the business at an abstract level. The logical data model offers a comprehensive formal structure to be used as a blueprint for business data and as input to define databases. And the physical data model specifies exactly how the logical data model will be implemented using a particular DBMS.

What Is Normalization?

In simple terms, [normalization](#) is the process of identifying the one best place where each fact belongs. Normalization is a design approach that minimizes data redundancy and optimizes data structures by systematically and properly placing data elements into the appropriate groupings. A normalized data model can be translated into a physical database that is organized correctly.

Normalization identifies the one best place where each fact belongs.

Normalization was created by Dr. E. F. Codd, the developer of the relational model, in the early 1970s. Like the relational model of data, normalization is based on the mathematical principles of set theory. Although normalization evolved from relational theory, the process of normalizing data is generally applicable to any type of data.

It is important to remember that normalization is a logical process and does not necessarily dictate physical database design. A normalized data model will ensure that each entity is well formed and that each attribute is assigned to the proper entity. Of course, the best situation is when a normalized logical data model can be physically implemented without major modifications. However, as we will learn in [Chapter 4](#), there are times when the physical database must differ from the logical data model due to physical implementation requirements and deficiencies in DBMS products.

The Normal Forms

As previously mentioned, Dr. Codd was the first to describe data normalization. Codd published several papers in 1971 and 1972 that described the first three normal forms. Subsequent work by Codd and others defined additional normal forms.

First Normal Form

The objective of *first normal form* (1NF) is to eliminate repeating groups and nonatomic data from an entity. When data conforms to 1NF, each attribute of the entity is a single discrete fact—in other words, atomic. The term *atomic* derives from *atom*, the smallest indivisible particle that can exist on its own.

First normal form eliminates repeating groups and non-atomic data from an entity.

Definition

A row is in first normal form if and only if all underlying domains contain atomic values only.

To normalize a data model into 1NF, eliminate repeating groups by splitting them into individual entities. In other words, do not use multiple attributes in a single entity to store similar data. Consider the sample data shown in [Table 3.2](#) for a STUDENT information system for a college or university.

Table 3.2. Unnormalized STUDENT Data

StudentID	StudentName	MajorID	StudentMajor	CourseNum	CourseName	CourseCompDate
2907	Smith, Jacob R	MAT	Mathematics	MAT0011	Discrete Math	2012-08-01
				MAT0027	Calculus I	2012-04-30
				EGL0010	English Classics I	2011-12-30
4019	Patterson, Jane K	PHI	Philosophy	PHI0010	Intro to Philosophy	2012-04-30
				CS00100	Programming Languages	2012-04-30
5145	Neeld, Norris B	EGL	English Literature	SOC0102	Ascent of Man	2012-08-01
6132	Morrison, Xavier Q	MUS	Music	MUS0002	Origin of Jazz	2012-04-30
				SOC0102	Ascent of Man	2012-08-01
7810	Brown, Richard E	CS	Computer Science			
8966	Juarez, Samantha	EGL	English Literature	EGL0010	English Classics I	2011-12-30
				EGL0101	Shakespeare II	2012-08-01

This data contains several violations of 1NF. First, we are tracking courses that really represent a repeating group for STUDENTs. So, the course information should be moved into a separate entity. Furthermore, we need to specify identifiers for both entities. The identifier is the primary key for the

entity.

Be careful to choose an appropriate primary key for each entity. This can be tricky. Your initial impulse may be to choose CourseNum for the COURSE entity. But in the case of this data, more information is required to identify the course information. The course completion date applies to a combination of STUDENT and COURSE—a course cannot be completed unless a student has enrolled in the course and is taking it.

A second violation of 1NF is the nonatomic data shown in the Student-Name attribute. A student name can be broken down into pieces: first name, middle initial, and last name. It is not indivisible and therefore violates first normal form. The end result in 1NF is depicted in [Tables 3.3](#) and [3.4](#).

Table 3.3. STUDENT Entity in 1NF

StudentID	LastName	FirstName	MiddleInit	MajorID	StudentMajor
2907	Smith	Jacob	R	MAT	Mathematics
4019	Patterson	Jane	K	PHI	Philosophy
5145	Neeld	Norris	B	EGL	English Literature
6132	Morrison	Xavier	Q	MUS	Music
7810	Brown	Richard	E	CS	Computer Science
8966	Juarez	Samantha		EGL	English Literature

Table 3.4. COURSE Entity in 1NF

StudentID	CourseNum	CourseName	CourseCompDate
2907	MAT0011	Discrete Math	2012-08-01
2907	MAT0027	Calculus I	2012-04-30
2907	EGL0010	English Classics I	2011-12-30
4019	PHI0010	Intro to Philosophy	2012-04-30
4019	CS00100	Programming Languages	2012-04-30
5145	SOC0102	Ascent of Man	2012-08-01
6132	MUS0002	Origin of Jazz	2012-04-30
6132	SOC0102	Ascent of Man	2012-08-01
8966	EGL0010	English Classics I	2011-12-30
8966	EGL0101	Shakespeare II	2012-08-01

Second Normal Form

Second normal form (2NF) ensures that all the attributes of each entity are dependent on the primary key. To transform 1NF data into 2NF, create separate entities for sets of attributes that apply to multiple records and assign a foreign key to the new entity to relate it to its previous entity. Simply stated, entity occurrences should not depend on anything other than the entity's primary key.

Second normal form ensures that all the attributes of each entity are dependent.

Definition

A row is in second normal form if and only if it is in first normal form and every non-key attribute is fully dependent on the key.

Let's once again turn our attention to [Tables 3.3](#) and [3.4](#). Notice that certain courses repeat in the COURSE entity, namely, "English Classics I" and "Ascent of Man." This situation indicates a violation of 2NF. To correct the problem, we need to identify the attributes that do not depend on the entire key and remove them. The removed attributes, along with the portion of the primary key on which they depend, are placed in a new entity, ENROLLMENT. The entire primary key of the original entity remains with the original entity.

Another benefit of the normalization process is that you will frequently encounter new attributes that need to be specified for the new entities that are created. For example, perhaps the new COURSE entity causes us to remember that each course is assigned a number of credits that count toward graduation. So, we create a new attribute to store the number of credits for each specific course. Of course, we may also decide we need more details on students such as address, phone number, and birth date, but we will omit these from our example to keep it simple.

The end results of normalization to 2NF are depicted in [Tables 3.3](#) (no changes were required for the STUDENT entity), [3.5](#), and [3.6](#).

Table 3.5. ENROLLMENT Entity in 2NF

StudentID	CourseNum	CourseCompDate
2907	MAT0011	2012-08-01
2907	MAT0027	2012-04-30
2907	EGL0010	2011-12-30
4019	PHI0010	2012-04-30
4019	CS00100	2012-04-30
5145	SOC0102	2012-08-01
6132	MUS0002	2012-04-30
6132	SOC0102	2012-08-01
8966	EGL0010	2011-12-30
8966	EGL0101	2012-08-01

Table 3.6. COURSE Entity in 2NF

CourseNum	CourseName	Credits
MAT0011	Discrete Math	3
MAT0027	Calculus I	4
EGL0010	English Classics I	3
PHI0010	Intro to Philosophy	3
CS00100	Programming Languages	3
SOC0102	Ascent of Man	3
MUS0002	Origin of Jazz	3

Third Normal Form

Third normal form (3NF) ensures that no relationships between attributes exist within an entity. Every attribute in the entity should depend only on the primary key. A tongue-in-cheek expression used to describe a data model in third normal is “Every attribute depends upon the key, the whole key, and nothing but the key, so help me Codd.”

Third normal form ensures that no relationships between attributes exist within an entity.

Definition

A row is in third normal form if and only if it is in second normal form and every non-key attribute is nontransitively dependent on the primary key.

A rule of thumb for identifying 3NF violations is to look for groups of attributes whose values can apply to more than a single entity occurrence. When you discover such attributes, move them to a separate entity.

It is time to review our STUDENT information again, this time looking for 3NF violations. Examine the STUDENT data in [Table 3.3](#) closely. Notice that students can have the same major and, as such, certain major information can be repeated; specifically, two students in our small sample are English Literature majors. To correct the problem, we need to remove major attributes that transitively depend on the key and create a new entity for them. [Tables 3.7](#) and [3.8](#) show the corrected data, now in 3NF.

Table 3.7. STUDENT Entity in 3NF

StudentID	LastName	FirstName	MiddleInit	MajorID
2907	Smith	Jacob	R	MAT
4019	Patterson	Jane	K	PHI
5145	Neeld	Norris	B	EGL
6132	Morrison	Xavier	Q	MUS
7810	Brown	Richard	E	CS
8966	Juarez	Samantha		EGL

Table 3.8. MAJOR Entity in 3NF

MajorID	StudentMajor
MAT	Mathematics
PHI	Philosophy
EGL	English Literature
MUS	Music
CS	Computer Science

A Normalized Data Model

To be complete, a diagram should be developed for the 3NF data model we just created for the STUDENT data. [Figure 3.7](#) shows such a data model. Notice that we have not filled in the optionality of the relationships. We could do this based on the sample data we used, but we really need to ask more questions before we can answer questions such as “Does every student have to have a major?” The current data shows this to be the case, but in reality, we know that most freshmen, and even upperclassmen, may attend college without having a formally declared major.

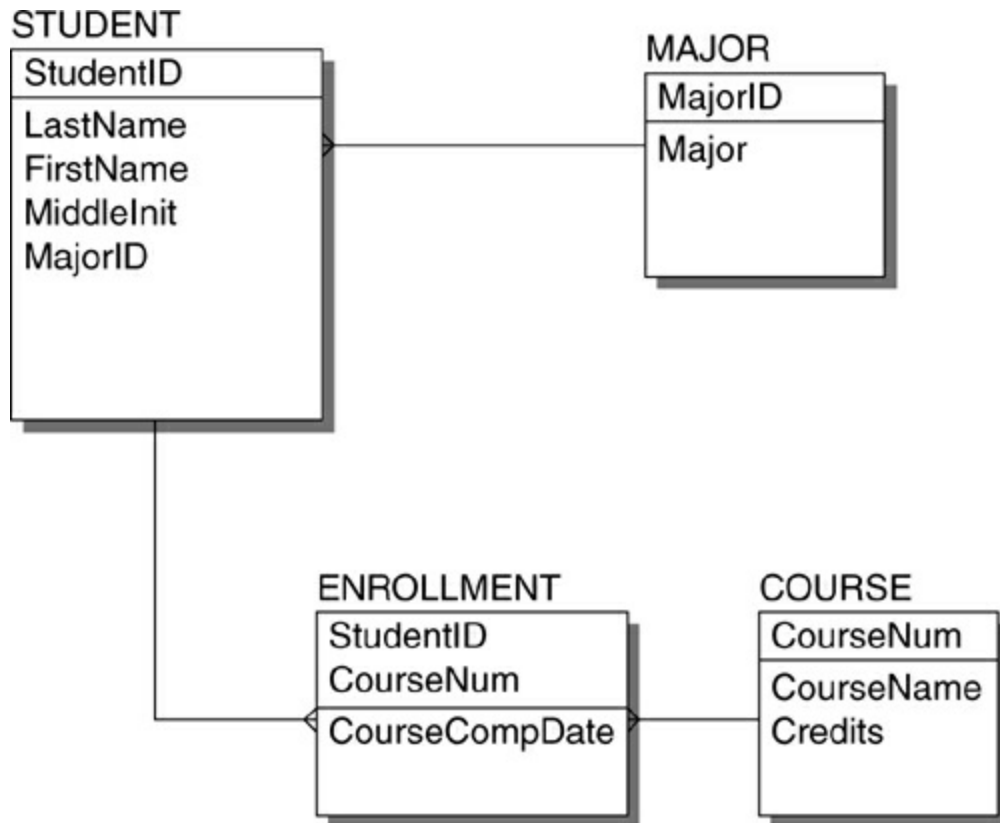


Figure 3.7. The STUDENT data model

Further Normal Forms

Normalization does not stop with 3NF. Additional normal forms have been identified and documented. However, normalization past 3NF does not occur often in normal practice. The following are additional normal forms:

Normalization past 3NF does not occur often in normal practice.

- *Boyce Codd normal form* (BCNF) is a further refinement of 3NF. Indeed, in his later writings Codd refers to BCNF as 3NF. A row is in Boyce Codd normal form if and only if every determinant is a candidate key. Most entities in 3NF are already in BCNF.
- *Fourth normal form* (4NF) states that no entity can have more than a single one-to-many relationship if the one-to-many attributes are independent of each other. An entity is in 4NF if and only if it is in 3NF and has no multiple sets of multivalued dependencies.
- *Fifth normal form* (5NF) specifies that every join dependency for the entity must be a consequence of its candidate keys.

Definition

A *determinant* is any attribute whose value determines other values within an entity occurrence.

Normalization in Practice

Our data models should be normalized as we move from the conceptual level to the logical level. By normalizing our data models, we ensure that we are creating a data blueprint for the best database design possible.

Normalization is a logical process.

But remember, normalization is a logical process. Many times data modelers try to mandate that the logical model become the physical database design with no changes. But it is not wise to force DBAs and technicians to forgo physical design. In the best of all worlds, a one-to-one logical-to-physical translation would work, but for many reasons this strict one-to-one mapping is not practical. We'll discuss the reasons for this in the next

chapter.

Relational Requires 1NF

The only normal form that is absolutely required for implementing a data model as a relational database is 1NF. 1NF ensures tabular data. Of course, further normalization to 3NF will make it easier to manage and maintain the data integrity for a relational database, but 3NF is not a hard-and-fast requirement for relational implementation.

Additional Data Modeling Issues

Please keep in mind that this chapter is intended as a broad overview of data modeling concepts for DBAs. Data modeling can be a full-time occupation for organizations that truly treat their data as a corporate asset. The information in this chapter should be sufficient for DBAs to work toward creating appropriate data models before creating databases. However, organizations that implement data administration will need to tackle more advanced data modeling issues such as the following:

- How should subtypes be handled? Subtypes can be created when a single entity has attributes that apply to only certain entity occurrences. For example, PART-TIME EMPLOYEE and FULL-TIME EMPLOYEE could be subtypes of the EMPLOYEE supertype entity. Subtypes may be needed if different relationships exist for the subtypes or if a sufficient number of attributes apply to only one of the subtypes.
- How is derived data to be treated for the purposes of a logical data model? If it is derived from other data, is it necessary to model it or document it?
- When should a domain become an entity? For example, business users in the United States typically know the abbreviations for states by heart. It is usually unnecessary to create an entity to map the state abbreviations to the state names. But not all such decisions are this simple.
- When time is a determining factor for data, how should it be handled in the data model? For example, perhaps a CUSTOMER can purchase only a single BASKET of goods at one time. But over time, the

CUSTOMER will purchase multiple BASKETs of goods; a new BASKET is purchased upon each visit.

- When should one-to-one relationships be allowed? Entities participating in a one-to-one relationship may be better collapsed into a single entity. But other issues such as relationships to other entities may impede this practice.

These and other issues will need to be studied and understood by data modelers and DAs in progressive, data-focused organizations.

Summary

Data modeling and normalization provide the foundation for database design. Data modeling is the practice of analyzing the things of interest to your organization and how these things are related to each other. Furthermore, the output is a data model that can be used as a blueprint for the data requirements of the business.

For data to be normalized, each entity must have a primary key for identification. All attributes of every entity must be atomic, indivisible, and not part of a repeating group. Furthermore, each entity contains attributes that apply to that entity only and depend on the full primary key for that entity only.

When data is normalized, facts are grouped into entities that make logical sense. No fact is associated with the *wrong* entity. The primary key wholly and completely identifies each attribute value of each entity occurrence. In other words, a normalized data model is a correct blueprint of the data.

The next phase is to implement a physical manifestation of the logical data model. To do so requires a thorough understanding of the DBMS you plan to use, as well as DBMS and database fundamentals.

Review

1. What is the difference between a conceptual data model and a logical data model?
2. Why is data modeling important for database development?
3. Describe the proper naming guidelines for entities and attributes.
4. Please provide a broad description of the goals of normalization.

5. What is the difference between an entity and an entity occurrence?
6. Review [Table 3.9](#) and explain why the data shown there violates 1NF.

Table 3.9. Sample Data

AccountNum	CustomerID	AccountTyp	BalanceAmt1	BalanceAmt2	BalanceAmt3
34109	105	A	2,352.47	5,123.31	
34134	113	A	15.95		
34147	422	A	18,549.22	100.00	1,799.85
34147	422	B	15.33		
34155	451	A	3,500.00	1.50	170.44
34162	990	X	500,000.00		
35010	802	X	1,908,190.12	10.88	
37891	999	Z	7.11		

7. Is the LastName attribute a good choice to be the primary key for the CUSTOMER entity? Please elaborate why or why not.
8. How is a relationship specified between two entities?
9. Every attribute will do one of three things. Name those things and describe them.
10. An attribute, or set of attributes, that identifies an entity occurrence is called what?

Bonus Question

Read the following paragraph and create a data model for the business under discussion.

“My partner and I own an independent record store. Well, I guess record store is really the wrong term these days since we really don’t sell a lot of records anymore. We sell mostly CDs, but also some tapes. And yes, we do stock vinyl records, too. Regardless of format (that is, CD, tape, or record), we sell long-playing albums as well as singles. Every item we sell is priced, and we keep track of the number of each individual title in stock on a daily basis. We like to track sales by both title and recording artist. The sales information by title needs to be tallied weekly to send to external ratings services so they can compile their best-seller charts. We often get customers who ask us to find certain songs for them. For singles, this is fairly easy because the title of the single is the title of the song. Some singles have multiple songs on them, though. Albums, on the other hand, comprise

multiple songs. We need the ability to search for titles by song to let customers know which titles (both singles and albums) contain the songs they are looking for.”

Suggested Reading

Allen, Sharon, and Evan Terry. *Beginning Relational Data Modeling*. 2nd ed. Berkeley, CA: Apress (2005). ISBN 1-59059-463-0.

Ambler, Scott W. *Agile Modeling*. New York, NY: John Wiley & Sons (2002). ISBN 0-471-20282-7.

Batini, Carlo, et al. *Conceptual Database Design: An Entity-Relationship Approach*. Redwood City, CA: Benjamin Cummings (1992). ISBN 0-8053-0244-1.

Bruce, Thomas A. *Designing Quality Databases with IDEF1X Information Models*. New York, NY: Dorset House (1992). ISBN 0-932633-18-8.

Carlis, John, and Joseph Maguire. *Mastering Data Modeling: A User-Driven Approach*. Boston, MA: Addison-Wesley (2001). ISBN 0-201-70045-X.

Chmura, Alan, and J. Mark Heumann. *Logical Data Modeling: What It Is and How to Do It*. New York: Springer Science + Business Media (2005). ISBN 0-387-22950-7.

Codd, E. F. “Further Normalization of the Database Relational Model.” In *Data Base Systems*, Courant Computer Science Symposia Series, vol. 6. Englewood Cliffs, NJ: Prentice Hall (1972).

Durell, William R. *The Complete Guide to Data Modeling*. Princeton, NJ: Data Administration, Inc. (1993). No ISBN.

Halpin, Terry, and Tony Morgan. *Information Modeling and Relational Databases*. 2nd ed. Burlington, MA: Morgan Kaufmann (2008). ISBN 978-0-12-373568-3.

Hay, David C. *Data Model Patterns*. New York, NY: Dorset House (1996). ISBN 0-932633-29-3.

Hoberman, Steve. *Data Modeling Made Simple*. Bradley Beach, NJ: Technics Publications (2005). ISBN 0-9771400-0-8.

Modell, Martin E. *Data Analysis, Data Modeling, and Classification*. New York, NY: McGraw-Hill (1992). ISBN 0-07-042634-1.

- Muller, Robert J. *Database Design for Smarties: Using UML for Data Modeling*. San Francisco, CA: Morgan Kaufmann (1999). ISBN 1-55860-515-0.
- Pilone, Dan, with Neil Pilman. *UML 2.0 in a Nutshell*. Sebastopol, CA: O'Reilly (2005). ISBN 0-596-00795-7.
- Ross, Ronald G. *Entity Modeling: Techniques and Application*. Boston, MA: Database Research Group (1988). ISBN 0-941049-00-0.
- Sanders, G. Lawrence. *Data Modeling*. Danvers, MA: Boyd & Fraser Publishing Company (1995). ISBN 0-87709-066-1.
- Schmidt, Bob. *Data Modeling for Information Professionals*. Upper Saddle River, NJ: Prentice Hall (1999). ISBN 0-13-080450-9.
- Silverston, Len, et al. *The Data Model Resource Book*. New York, NY: John Wiley & Sons (2002). ISBN 0-471-15264-8.
- Simsion, Graeme C., and Graham C. Witt. *Data Modeling Essentials*. 3rd ed. San Francisco, CA: Morgan Kaufmann (2005). ISBN 0-12-644551-6.
- Weilkiens, Tim, and Bernd Oestereich. *UML 2 Certification Guide*. San Francisco, CA: Morgan Kaufmann (2007). ISBN 978-0-12-373585-0.

4. Database Design

Database design is the process of transforming a logical data model into a physical database design and then implementing the physical model as an actual database. More precisely, database design requires up-front data modeling and normalization as discussed in [Chapter 3](#). A logical data model is required before you can even begin to design a physical database. This chapter assumes that the logical data model is complete. The focus of discussion, therefore, will be on producing a physical database from the logical data model.

From Logical Model to Physical Database

The physical data model is created by transforming the logical data model into a physical implementation based on the DBMS to be used for deployment. To successfully create a physical database design you will need to have a good working knowledge of the features of the DBMS, including

- In-depth knowledge of the database objects supported by the DBMS and the physical structures and files required to support those objects
- Details regarding the manner in which the DBMS supports indexing, referential integrity, constraints, data types, and other features that augment the functionality of database objects
- Detailed knowledge of new and obsolete features for particular versions or releases of the DBMS
- Knowledge of the DBMS configuration parameters that are in place
- Data Definition Language (DDL) skills to translate the physical design into actual database objects

The physical data model is created by transforming the logical data model into a physical implementation.

Armed with the correct information, you can create an effective and efficient database from a logical data model. The first step in transforming a logical data model into a physical model is to perform a simple translation from logical terms to physical objects. Of course, this simple transformation will not result in a complete and correct physical database design—it is

simply the first step. Let's address the transformation step by step.

Transform Entities to Tables

The physical counterpart of an entity is a table. Therefore, the first step in transforming a logical data model into a physical database is to map each entity in the data model to a table in the database. The final database that you implement need not adhere to this strict one-to-one entity-to-table mapping. For example, you may need to consolidate or break apart tables—a process called denormalizing—for performance reasons. Denormalization will be covered in detail later in this chapter.

Map each entity in the data model to a table in the database.

In general, though, do not initially deviate from the simple rule of creating a table for each entity in the logical data model. The logical data model represents the “things” of interest to the business. During the data modeling process, these things were analyzed and designed as entities, each having a specific identity and purpose for existence. The only reason to deviate from this process is if application performance or data availability would be unacceptable without a change.

Transform Attributes to Columns

The physical counterpart of an attribute is a column within a table. When you map entities to tables, map the *attributes* of each entity to the *columns* of each respective table. At least initially, do not change the basic definition of the columns. For example, do not group attributes together into a composite column.

The attributes of each entity should be mapped to the columns of each respective table.

Try to maintain the same naming convention for physical columns as was used for logical attribute names. However, you must understand that the physical constraints of the DBMS being used may limit your ability to do so. Always take into account the capabilities and limitations of the DBMS when creating the physical database from the logical data model.

Transform Domains to Data Types

To support the mapping of attributes to table columns you will need to map

each logical domain of the attribute to a physical data type, perhaps coupled with additional constraints. Each column must be assigned a [data type](#). Certain data types require you to specify a maximum length. For example, you could specify a character data type as CHAR(20), indicating that up to 20 characters can be stored for the column. You may need to apply a length to other data types as well, such as graphic, floating-point, and decimal (which also require a scale).

Map each logical domain to a physical data type, perhaps coupled with additional constraints.

Commercial DBMS products do not support domains, so the domain assigned in the logical data model must be mapped to a data type supported by the DBMS. You may need to adjust the data type based on the DBMS you use. For example, what data type and length will be used for monetary values if no built-in currency data type exists? Many of the major DBMS products support user-defined data types, so you might want to consider creating a data type to support the logical domain if no built-in data type is acceptable.

There may be multiple physical data types that can be used successfully for a domain. Consider a logical domain whose valid values are integers between 1 and 10. You could choose an integer data type, of which there may be several (e.g., BIGINT, INTEGER, SMALLINT, TINYINT). Alternatively, you could choose a decimal data type with a zero scale. You might even choose to store the data in a 2-byte character column if no mathematical operations are required. Any of these can work. As the DBA, you will need to determine which data type can be most efficiently accessed, stored, maintained, and processed by the applications accessing the data. To make such a decision requires in-depth technical knowledge of the way in which your DBMS physically stores each type of data, as well as knowledge of application processing details. Data typing should not be undertaken hastily. Choosing an inappropriate data type could have ramifications in terms of both storage space and disk I/O.

In addition to a data type and length, you may also need to apply a [constraint](#) to the column. Consider, once again, the domain of integers 1 through 10. Simply assigning the physical column to an integer data type is insufficient to match the domain. A constraint must be added to restrict the values that can be stored for the column to the specified range of 1 through

10. Without a constraint, negative numbers, zero, and values greater than 10 could be stored. Using check constraints, you can place limits on the data values that can be stored in a column or set of columns. Check constraints are covered in detail in [Chapter 13, “Data Integrity.”](#)

The nullability of each column in the table must also be specified. The logical data model should contain information on the nullability of each attribute, and this information can be copied for each requisite column in the physical database. Some DBMS software enables you to assign a default value to be used when a row is to be inserted and no value has been provided for the column. You must assign column default values when you create the table.

Specify the nullability of each column.

For text or character data, you need to make an additional decision: Should the column be fixed length or variable length? A *fixed-length column* occupies a preset and unchanging amount of storage for each row. A *variable-length column* specifies a maximum size, but the actual length used by the column can vary for each row. Variable-length columns can save storage space but usually require additional manipulation by application code to insert and maintain them. Another negative aspect is that variable-length columns can cause the table to require more frequent attention. Changes made to row size will have the effect of moving rows within the database. If the DBMS offers a compression option, you may be able to save more space by compressing the database than by implementing variable-length columns.

Primary Keys

Specification of a primary key is an integral part of the physical design of entities and attributes. When designing the logical data model, you assigned a primary key for every entity, and as a first course of action you should try to use that primary key. However, multiple candidate keys are often uncovered during the data modeling process. For physical implementation you may decide to choose a primary key other than the one selected during logical design—either one of the candidate keys or another surrogate key.

If the primary key is unwieldy, you may need to choose another primary key. Perhaps the key is composed of several columns or is a nonstandard data type. In either of these cases it may be better to choose a surrogate primary

key. Most DBMS products provide built-in features that can assist in the definition of primary keys. Some examples include support for default values, ROWID data types, SEQUENCE objects, and the identity property (see the sidebar).

Most DBMS products provide built-in features to help define primary keys.

As a rule of thumb, though, be sure to identify a primary key for each physical table you create. Failure to do so will make processing the data in that table more difficult. If there is no key to uniquely identify rows of the table, it will be difficult for programs to specifically select, modify, and delete individual rows. Furthermore, without a primary key, dependent relationships cannot be defined for that table.

Be sure to identify a primary key for each physical table you create.

There are, of course, exceptions to the general rule. For example, a table implemented for logging or that serves as a generic queue may not require a primary key. Before creating any table without a primary key, though, be sure that you have a solid reason for deliberately defining a table without a unique primary key.

The Identity Property

The [*identity property*](#) is a feature supported by several of the most popular relational DBMS products. It can be assigned to a column that has a numeric (usually integer) data type. When the identity property is assigned to a column, the DBMS treats that column in a special way. The database user does not provide values for the column when rows are inserted into the table in which the column exists. Instead, the DBMS increments a counter and automatically uses that value for the column. Usually only one column per table can be assigned the identity property.

The identity property provides an efficient way for ever-increasing sequential values to be populated into a relational table. Before using this feature, make sure that your DBMS supports the identity property and that you completely understand how the feature works and its impact on functionality and performance. For example, does the DBMS

provide a way to reset the identity property counter on a table-by-table basis?

Column Ordering

Before implementing a physical table, be sure to review the order of the columns. The order in which columns are specified is irrelevant from an operational perspective—that is, the DBMS will produce the same results regardless of the sequence of the columns in the table. Nevertheless, the efficiency of how those results are obtained can vary greatly. Column sequencing can impact performance; therefore, for physical implementation you may need to change the sequence recorded in the logical data model.

Let's take a closer look. The way in which the DBMS logs changes can impact performance. DB2, for example, logs database modifications from the first byte changed to the last byte changed. The exception is variable-length rows, in which case DB2 logs a change from the first byte changed to the end of the row.

So, to take advantage of this knowledge about DB2's physical implementation, we should sequence the columns based on how DB2 logs. Infrequently updated nonvariable columns should be grouped together at the beginning of the table, followed by static (infrequently updated) variable columns and then frequently updated columns. This structure will ensure that the least amount of data required is logged, thereby speeding up any data modification processes. Another good idea would be to group together any columns that are frequently modified. This can also reduce the amount of data logged. Because each DBMS logs data differently, you will need to understand how your DBMS logs and how column ordering can impact performance.

Build Referential Constraints for All Relationships

The physical counterpart of a relationship is a [referential constraint](#). To define a referential constraint you must create a primary key in the parent table and a foreign key in the dependent table. The referential constraint ties the primary key to the foreign key.

The referential constraint ties the primary key to the foreign key.

Referential Integrity

It is not sufficient merely to identify primary and foreign keys that make up relationships between tables. The functionality of each relationship is greatly affected by the parameters chosen for the referential constraint and the values in the foreign key column(s). A set of rules, which is applied to each relationship, determines the status of foreign key columns when inserted or updated, and of dependent rows when a primary key row is deleted. For example, when a primary key is deleted that refers to existing foreign key values, the rule specifies whether the DBMS should void the primary key deletion, delete the foreign key values too, or set the foreign key values to null.

In general, a foreign key should either contain a value within the domain of foreign key values or be null. Any other value is unacceptable, and the referential constraint will cause the DBMS to reject such values during operation. *Referential integrity*, or RI for short, guarantees that an acceptable value is always in each foreign key column.

Referential integrity guarantees that an acceptable value is always in each foreign key column.

It is a good physical design practice to implement referential integrity using database constraints instead of trying to program integrity into application programs. Using database RI will ensure that integrity is maintained whether data is changed in a planned manner through an application program or in an ad hoc manner through SQL statements or query tools.

Of course, there are exceptions to every rule. Cases where you should avoid using database-enforced RI constraints include code tables where the codes do not change frequently and data warehouse implementations where the data is propagated from a referentially intact source. Basically, do not use RI as a substitute for performing data validation and edit checks. Listing valid values in a check constraint can be a better solution than RI for static data with a small number of valid values.

[Chapter 13, “Data Integrity,”](#) discusses referential integrity in more detail.

Build Physical Data Structures

Designing and implementing a physical database from a logical data model is

not just a simple matter of mapping entities to tables, attributes to columns, and relationships to referential constraints. Quite a few other database design issues must be addressed. One of these issues is preparing for the physical storage of data.

Although relational data is expressed to the user by means of a table, underlying files or data sets must exist to store the actual data—and those files are not necessarily stored as a simple grid of rows and columns. During the physical design process, the DBA must map each table to a physical structure to store the table's data. These physical structures are commonly called *tablespaces* (or *data spaces*). As shown in [Figure 4.1](#), a database comprises one or more tablespaces, and each tablespace contains one or more tables. Depending on the DBMS, a table may be able to span multiple tablespaces, too. The DBA decides how to map tables to tablespaces based on the anticipated usage of the data, the type of tablespace, and the features of the DBMS. Please see the sidebar “[SQL Server Filegroups](#)” for information on a similar file structure used by Microsoft SQL Server.

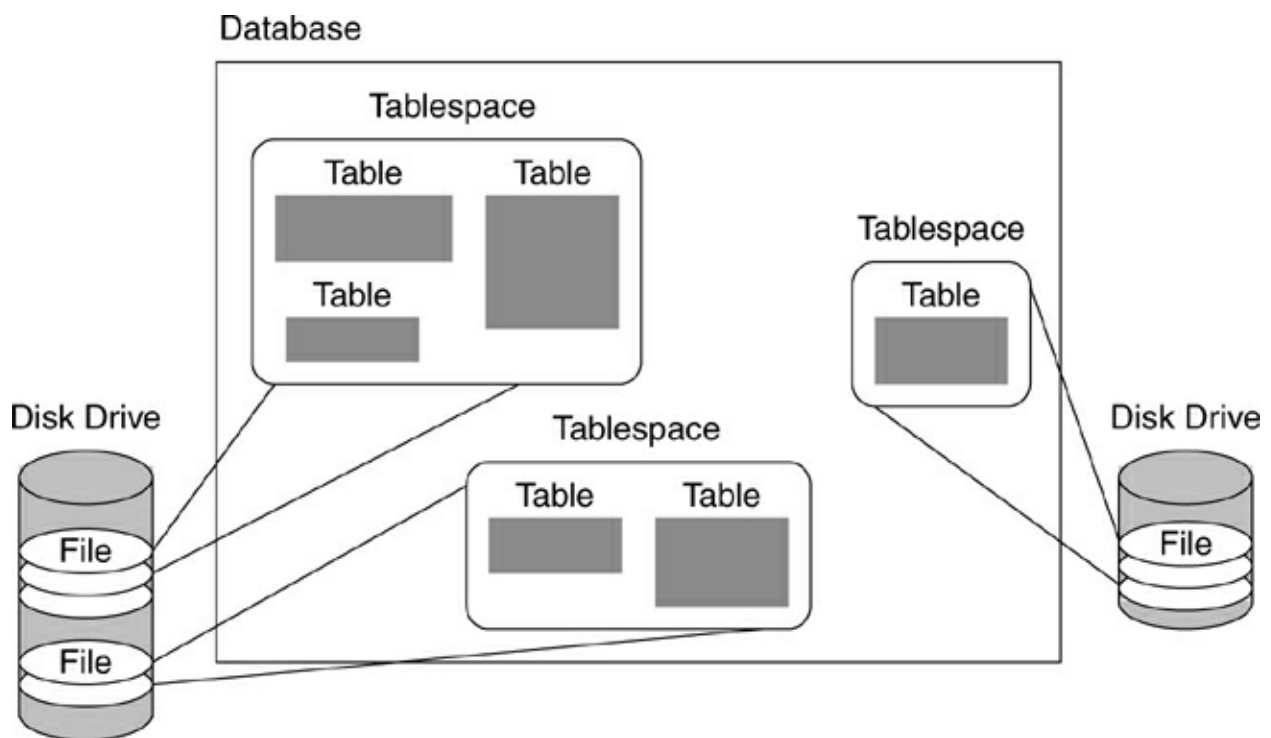


Figure 4.1. Mapping files to database structures

An additional physical design process is the effective planning for storage and space usage. To calculate the amount of storage required to store a table, the DBA must first establish the *row size*. This is accomplished by adding up

the maximum size of all the columns, based on their assigned data type and length. An average size can be substituted for variable-length columns. An estimate for the number of bytes required can be calculated by multiplying the row size by the number of rows planned to be stored in the table. Of course, the DBA must also factor in any storage overhead required by the DBMS for row and page headers, pointers, and the like. Because each DBMS uses different techniques, each DBMS will have different overhead requirements.

SQL Server Filegroups

Microsoft SQL Server provides filegroups to group together the operating system files containing data from a single database. By using a filegroup, the DBA can simplify certain administration tasks. A filegroup is a property of a SQL Server database. A filegroup cannot contain the operating system files of more than one database, though a single database can contain more than one filegroup.

Every database is created on a single filegroup named PRIMARY. After the database is created, additional filegroups can be added. If your database is very large and active, multiple files can be used to improve performance.

When using SQL Server with very large databases (VLDBs), filegroups are useful for dividing the database into components for backup and restore. It can be difficult to manage VLDB backups and restores without using filegroups.

In a nutshell, SQL Server filegroups can be used to provide greater control over the I/O of a given database.

The DBA must also determine the type of files to be used, based on the operating system being used as the database server. In some cases, database files that do not use the file system of the operating system can be more efficient. Such files are sometimes referred to as *raw files*. Because a raw file can prevent operating system overhead, it is often a better choice for database implementation. The trade-off is that it requires more administrative support from the DBA.

To prepare storage devices for the database, though, the DBA must

determine the size of not only table structures but also index structures. Indexes are covered in more depth in the next section.

The DBA must determine the size of not only table structures but also index structures.

The DBA will also build some free space into the database design. When free space is assigned within the database, the DBMS can add data more efficiently and with fewer outages for reorganizing and restructuring.

Solid-state devices can be helpful, too. Storing databases that have extreme performance requirements on solid-state devices instead of disk can improve the performance of many database operations.

A final storage consideration is whether to use *compression*. Most DBMS products provide methods for compressing data. If the DBMS provides no compression option, third-party compression products can be purchased. With compression, less storage is required to store the same amount of data. A compression routine will algorithmically compress data as it is added to the database and decompress the data as it is retrieved. Compression adds CPU overhead¹ to database processing because access requires the DBMS to compress or decompress the data. However, I/O may be more efficient with compression: Because the data takes up less space on disk, more data can be read with a single I/O operation.

Database Performance Design

When implementing a physical database from a logical data model, you must begin to consider how the database will perform when applications make requests to access and modify data. A basic fact of database processing is that disk access is slower than memory access—slower by orders of magnitude. If the DBMS were required, in every instance, to scan through the database row by row, or block by block, looking for the requested data, no one could afford to use databases. Fortunately, several good techniques exist to allow data in the database to be accessed more rapidly.

Designing Indexes

One of the best techniques for achieving acceptable query performance is the creation of appropriate indexes on your database tables. Of course, the trick is in determining how many indexes to create and how exactly to define each

index. First, let's cover some index basics.

An index is an alternate path to data in the database. The structure of an index makes it easier to find data in the database, with fewer I/O operations. Therefore, queries can perform faster when using an index to look up data based on specific key values.

Many types of indexes are supported by the major DBMS products. Indexes can be unique or nonunique, clustering or nonclustering, single column or multicolumn. An index can be structured as a b-tree or a bitmap. Some DBMS products even support hashing indexes. However, the basic goal of every index is to optimize query processing.

An index is an alternate path to data in the database.

In a relational system, the DBMS—not the programmer—decides whether to use an index. Therefore, the DBA must design indexes based on the types of queries that will be run against the database. The DBA must understand the operation of the relational optimizer and design indexes that are likely to be used during application processing. In the absence of an appropriate index for a query, the DBMS will most likely revert to a *table scan*—every row of the table will be read to determine whether the data matches the query specifications. Table scans are costly if the intent of the query is not to process every row in the table.

In general, try to build indexes on large tables to support the most frequently run queries. Queries that access 25 percent or fewer of the table's rows are good candidates for indexing. When more than 25 percent of the table's rows are to be selected, you should determine whether an index would be more efficient than a table scan. This will differ from query to query, and from DBMS to DBMS. Use the tools provided by the DBMS to determine the effectiveness of your indexes. (The SHOWPLAN or EXPLAIN command will show whether an index is used for a particular query.)

Furthermore, you should create indexes on the most-referenced columns in frequently run queries in your application. The order in which columns appear in an index is important. By choosing the right order, you may be able to make a particular index available to many other queries. For example, if quite a few application queries look for items based on ItemType and a few other queries look for items by ItemType and ItemColor, a single composite index on the combination of (ItemType, ItemColor) can satisfy the needs of

both types of queries.

Create indexes on the most-referenced columns in frequently run queries in your application.

A single table can have multiple indexes defined for it; you are not limited to a single index. Plan on creating indexes for the most common access requirements in each application's SQL WHERE and JOIN clauses. The following situations should prompt you to consider creating an index:

- *Foreign keys*. Even if the DBMS does not require foreign key columns to be indexed, it is a good idea to do so. Creating an index on foreign key columns can enhance the performance of joins based on the relationships between the tables, but it may also speed up the DBMS's internal processing to enforce referential integrity.
- *Primary keys*. An index is usually required on the primary key columns to enforce uniqueness.
- *Candidate keys*. Even though indexes are not required on candidate keys, it is a good idea to index them if processes will look up data based on the candidate key.
- *Index-only access*. If all of the columns in a data retrieval request exist in an index, it may be possible to satisfy the request using only the index. Avoiding I/O to the table can enhance performance. Therefore, it is sometimes a good idea to overload an index with columns to facilitate index-only access for certain requests. Such an index is sometimes called a "covering index" because the index "covers" all of the data required by certain queries. Additionally, some DBMS products provide the capability to add columns to a unique index without having the column enforced for uniqueness. This technique can be used to pack additional columns into a single index instead of requiring multiple indexes.
- *Sorting*. Another reason for building indexes is to minimize sorting. Queries that use JOINS, ORDER BY, GROUP BY, UNION, and DISTINCT can cause the DBMS to sort the intermediate or final results of that query. If indexes are created to support these features, the DBMS may be able to use the index for ordering and avoid invoking a costly sort.

A frequent error made by DBAs and performance analysts is to index by object instead of indexing by workload. Indexing by object means creating indexes as you create your tables. Oftentimes, during the physical database implementation process, a DBA will create a database, then groups of tablespaces and tables. And every time a new table is created, the DBA will attempt to create the indexes on that table. This approach is not optimal.

Instead, DBAs should build indexes based on workload. Indexes should support the predicates in the SQL that is written to access your tables. Building indexes to support predicates of the most frequently executed queries and most important queries should be your first indexing step after building the unique indexes required to support primary keys and unique constraints.

Of course, this requires knowledge of how your tables will be accessed. And when you are first creating tables, you will not have any SQL. Sometimes you may have vague pseudo code descriptions of potential queries, but you won't have an accurate picture of access. Therefore, indexing has to be an incremental task, performed on an ongoing basis as code is written against your databases.

As you continually monitor and build new indexes, be sure to review the old ones that were created. Sometimes a new index can render an existing index obsolete. It is a good idea to drop indexes that are not used for query optimization because each index must be maintained as data is modified. This can negatively impact database and application performance.

Exercise great care in the creation of an indexing scheme for your databases. Be sure to analyze all of the data access requirements of your applications to ensure optimal indexing. You simply cannot design proper indexes without knowledge of how tables are to be accessed. Furthermore, you need a comprehensive view of the access requirements. An index that is the best solution for a particular query can potentially adversely affect the performance of other queries in the system.

Exercise great care in the creation of an indexing scheme for your databases.

Additionally, be aware that indexes do not come for free. The DBMS must keep the indexes updated as the table's data is inserted, updated, and deleted. For that reason, you may want to avoid building indexes on columns that are

frequently modified, if you can. Make sure that every index you create provides benefit to the performance of a query or set of queries without significantly degrading the overall performance of the applications accessing the database. Here are some things to consider when determining the cost of an index:

- Additional overhead is incurred to update the index when rows are inserted and deleted, or when indexed columns are updated in the base table.
- Additional disk space is required to store indexes. Some DBMS products offer index compression as an option. For very large databases, index compression can be worthwhile.
- Utilities such as LOAD and REORG may take longer to run against a table with many indexes because the indexes must also be maintained during the utility processing.
- Additional files are required to store the indexes, which could potentially cause operating system problems if the number of files that can be open at one time is exceeded.

When building indexes, keep in mind that they will be used in conjunction with the base table data. Consider allocating separate DBMS buffer areas for caching index reads instead of using the same buffers used for data. Index entries will be smaller than a full table row, and thus more index entries can be kept in memory longer if they are not combined with table data in the database buffers. You might also consider placing the indexes and the table data on different disk drives to minimize disk seek time.

DBAs must learn how each DBMS supports indexing so they can create indexes to support the data access requirements of their database applications. Let's learn a little bit more about the types of indexes that your DBMS may support: b-tree, bitmap, reverse key, partitioned, and ordered.

B-Tree Indexes

The basic indexing technique supported by most relational database systems is the b-tree index. A [*b-tree*](#) is a keyed, treelike index structure. A b-tree index begins at the root page and fans out down to the leaf pages. [Figure 4.2](#) shows the basic structure of a b-tree index.

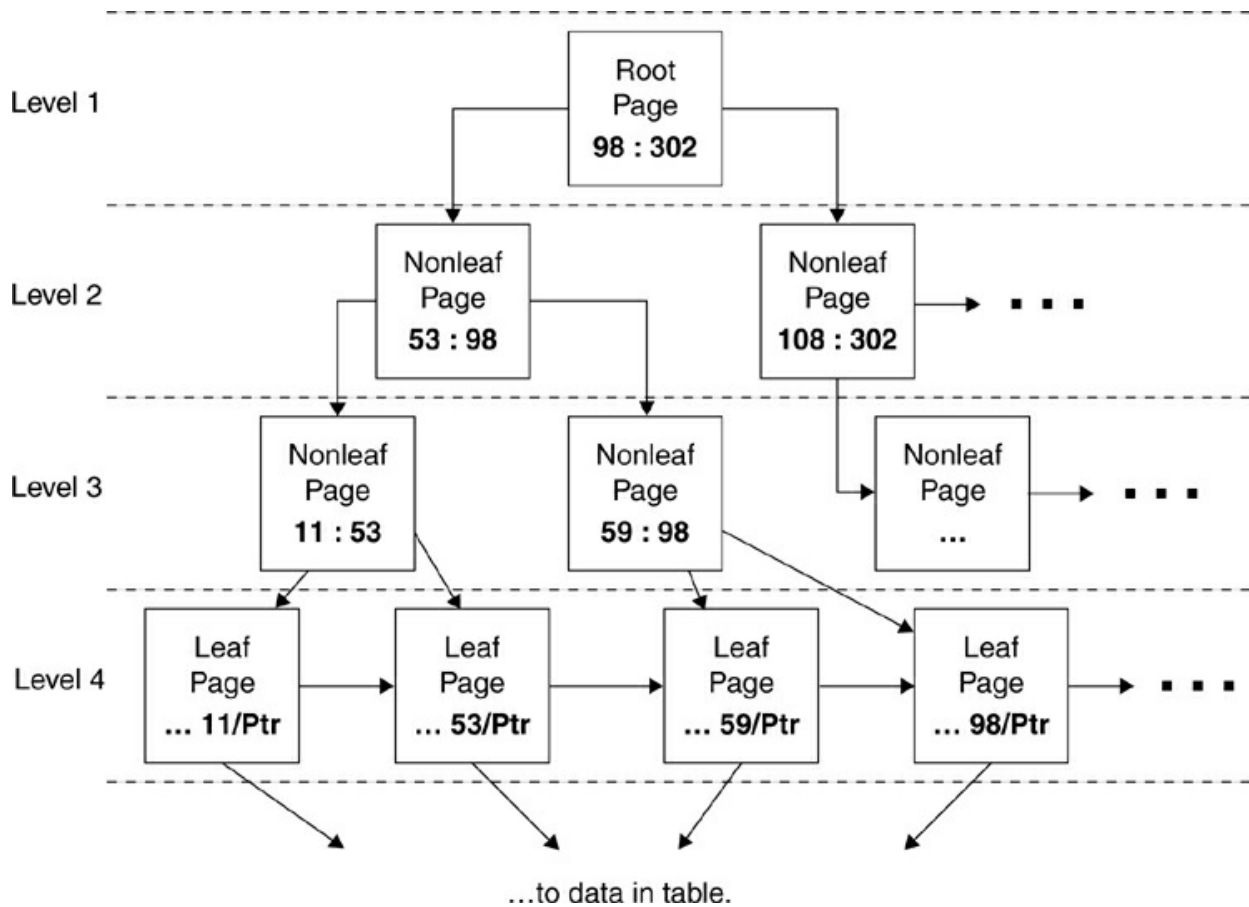


Figure 4.2. B-tree index structure

A b-tree is a keyed, treelike index structure.

The pages of a b-tree index are referred to as *nodes*. Nodes exist in levels in the b-tree, with nodes above the leaf level containing directory entries and pointers to lower-level b-tree nodes. Nodes at the lowest level are called *leaf pages*. Leaf pages contain entries with the key value and a pointer to individual data rows in the table. As data is added to the table, the b-tree index is updated by storing the new key values in the appropriate location in the index structure. To access data using the index, the DBMS begins at the root page and follows the pointers through the index until the key value is located at the leaf level, where a pointer leads to the actual table data. Each parent node contains the highest key value that is stored in its direct dependent nodes. The leaf pages of a b-tree index can be scanned for ranges of values once the key has been looked up.

Refer to [Figure 4.2](#) again. Suppose a query is run with a search condition

requesting data where the key value is equal to 53. The DBMS can traverse the index and in this case will wind up on the second leaf page. This leaf page contains a pointer to the actual row in the table containing the requested key value. A maximum of five I/O requests is required to satisfy this query: one for a page at each level of the index and an additional request for the table page.

An access request using an index can perform better than a table scan because the requested data can be accessed directly using the pointers from the leaf node of the index. This reduces I/O and enhances performance for most data access requests.

Bitmap Indexes

Bitmap indexes are a different type of index supported by some DBMS products. A bitmap index solves only a narrow range of problems but provides superb performance. A bitmap index is most useful for query-heavy tables that are infrequently modified. Furthermore, bitmap indexes are most useful where the columns to be indexed have a very small number of distinct values, such as sex (Male/Female) or Boolean (True/False) data. Data warehouses and data marts can often benefit from bitmap indexes.

A bitmap index solves only a narrow range of problems but provides superb performance.

A bitmap index is really quite simple. The implementation of a bitmap index is, appropriately enough, accomplished using a string of zeroes and ones, or bits. For each key value of the bitmap index a separate string of zeroes and ones is stored. The number of distinct key values determines the number of bit strings. For example, a bitmap index defined for the Sex column of the EMPLOYEE table might have three strings, one for male, one for female, and one for unknown. A bitmap index on the State column (for states within the United States) could have 51 strings—one for each state and an extra one for unknown.

To elaborate, let's consider the Sex example again. A bitmap index is created on the Sex column of an EMPLOYEE table that contains ten rows. The bitmap index will contain three strings as indicated above, each with 10 bits. The string is positional. Whatever position has a bit turned on ("1"), the Sex column in that row will contain the value for which that particular string

was built. Examine the following three bitmaps:

```
'Male' 1000011101  
'Female' 0110000010  
'Unknown' 0001100000
```

These strings indicate that rows 1, 6, 7, 8, and 10 are males; rows 2, 3, and 9 are females; and rows 4 and 5 are unknown.

Finding the set of records with any of several values that are bitmap indexed simply requires adding the strings for those values. Bitmaps can be much faster than table scans and even b-tree index retrieval under the right circumstances. The strings of the bitmap index can be small enough to maintain in memory, minimizing I/O operations. Even for queries that retrieve large numbers of rows, a complex query using bitmaps can be very efficient.

The problem with bitmaps is that a separate string is needed for each value occurring in a field. When a column can take on a large number of different values, a bitmap index is not practical because too many strings will be required. Additionally, the zeroes and ones of the bitmap cannot be used for calculations or be read directly to determine the actual values they represent. Such limitations make true bitmap indexes impractical for most applications. Some DBMS products have extended the bitmap index to make it more practical for columns with higher cardinality. If your DBMS supports bitmap indexes, make sure you understand the exact nature of the bitmaps being used and the situations under which bitmap indexes are practical.

Reverse Key Indexes

A *reverse key index* is basically a b-tree index where the order of bytes of each indexed column is reversed. The order of the columns within the index is not reversed, just the bytes within each column. Such indexes can be useful to eliminate data hot spots in OLTP applications. When the byte order is reversed, adjacent key values are not physically stored together. So reverse key indexes help to distribute the otherwise concentrated index data across leaf nodes, thereby improving performance.

Reverse key indexes can eliminate data hot spots in OLTP applications.

So, if a reverse key index is created on the FirstName column of the

EMPLOYEE table, and the value “Craig” is inserted, the value “giarC” is used instead.

If the DBMS you are using does not support reverse key indexes, you might be able to programmatically duplicate the effect. To do so, you will need to use program logic (or perhaps an exit routine, if supported by the DBMS) to reverse the values before inserting them into the columns. Of course, you will need to programmatically unreverse the values when they are retrieved. This process does not work well if the data must be queried in an ad hoc manner, outside the scope of a program.

Partitioned Indexes

Partitioned indexes are basically b-tree indexes that specify how to break up the index (and perhaps the underlying table) into separate chunks, or partitions. Partitioning is usually done to enhance performance and increase availability. When data is spread out across multiple partitions, you may be able to operate on one partition without impacting others—for example, to run utilities, to take data offline, or to place the underlying files on separate disks.

Partitioning is usually done to enhance performance.

Most DBMS products support partitioning, but in different ways. Be sure to understand the nuances of your particular DBMS implementation before partitioning.

Ordered Indexes

Most DBMS products provide an option to specify the order in which b-tree key values are ordered. The order specified, either ascending or descending, will impact the usability of an index to avoid sort operations or to minimize I/O requirements for retrieving MIN or MAX values. Create indexes in the appropriate order to support the types of queries that are to be run against the table.

Index Summary

Indexing is an important component of a physical database design. Indeed, the single most important thing a DBA can do to optimize the performance of database applications is to create effective indexes. In order to do so, the DBA needs to know what indexing options are available in the DBMS being

used, but more importantly, the DBA must be able to match the DBMS indexing options to the type of processing to be performed against the table. Only by examining the SQL statements that operate on the database tables can an effective indexing strategy be developed. Finally, keep in mind that special processing requirements may require special indexing needs and that add-on products are available from ISVs that can augment the indexing options available to you.

Hashing

Hashing is a technique that uses key values to enable quick direct access to data. An algorithm is used to transform the key values into a pointer to the physical location of the rows that have those key values. The algorithm is typically referred to as a randomizer, because the goal of the hashing routine is to spread the key values evenly throughout the physical storage.

Hashing uses key values to enable quick direct access to data.

In general, the better the randomizing algorithm, the better the results of hashing will be. When the randomizer generates the same pointer for two different key values, a collision occurs. Different techniques can be used to resolve collisions. Typically, the collision resolution algorithm attempts to keep the data on the same page to avoid additional I/O. When pages fill up and collisions force the data to another page, performance rapidly degrades.

Hashing works better with a large amount of free space. A disadvantage of hashing is the amount of space that must be preallocated for data to be hashed into.

Hashing has the big advantage that normally only one database I/O request is needed to retrieve a row of data using the key. Hashing works best for direct data lookup of one row or a small number of rows. If you need to retrieve ranges of data, hashing is not optimal because the data will be spread out instead of clustered and therefore I/O costs will be substantial. Additionally, hashing requires a unique key to minimize collisions.

So, hashing should be considered only when an overwhelming majority of the queries against the table are based on lookups using the key and will return small results sets.

Clustering

Clustering describes a way of storing table data physically. The term refers to keeping rows in a specific order on the disk. Through clustering, data that is commonly accessed together can be stored together on the same or contiguous database pages. Clustering optimizes performance because fewer I/O requests are required to retrieve data.

Clustering optimizes performance because fewer I/O requests are required to retrieve data.

Actually, more accurately, for some products clustering indicates that the DBMS should *attempt* to maintain rows in the sequence of specific column values. If insufficient space is available to maintain clustering when data is inserted or modified, the DBMS typically stores the data without forcing clustering. Therefore, a clustered table may not actually be 100 percent clustered by the key value at all times.

Usually an index, called a clustering index, is required to support clustering. The columns identified as the index key indicate how the table upon which the index is defined should be clustered.

Consider clustering tables under the following circumstances:

- When a large number of queries retrieve ranges of data based on specific column values.
- When a foreign key exists in the table. A foreign key typically represents the “many” end of a one-to-many relationship. It is common for queries to request data by the foreign key, resulting in large sequential reads.
- When data is frequently sorted together (ORDER BY, GROUP BY, UNION, SELECT DISTINCT, JOINS).

When you cluster a table, be sure to consider the frequency of modification. Inserts and updates can cause data to become unclustered. Favor clustering infrequently modified data over very frequently modified data. However, the primary key is almost always a bad choice for clustering because primary key access tends to be random and clustering optimizes sequential access.

Be sure to understand how your DBMS implements clustering. Some DBMS products merge the table and the clustering index into a single structure, which may require you to modify your administration techniques

and procedures. Other clustering differences exist, too. For example, although Oracle supports a structure called a cluster, it does not perform clustering as just described; instead, it interleaves index keys for multiple tables.

Interleaving Data

When data from two tables is frequently joined, it can make sense to use the join criteria to physically interleave the data into the same physical storage structure. *Interleaving* can be viewed as a specialized form of clustering.

To better understand data interleaving, refer to [Figure 4.3](#). The dots indicate rows in different tables. The data is interleaved on the disk, based on the join criteria. Notice that the light dots (table 1) are intermixed in the same file as the dark dots (table 2). When data is interleaved this way, join performance may improve—but only for the specific join that the data was interleaved to optimize.

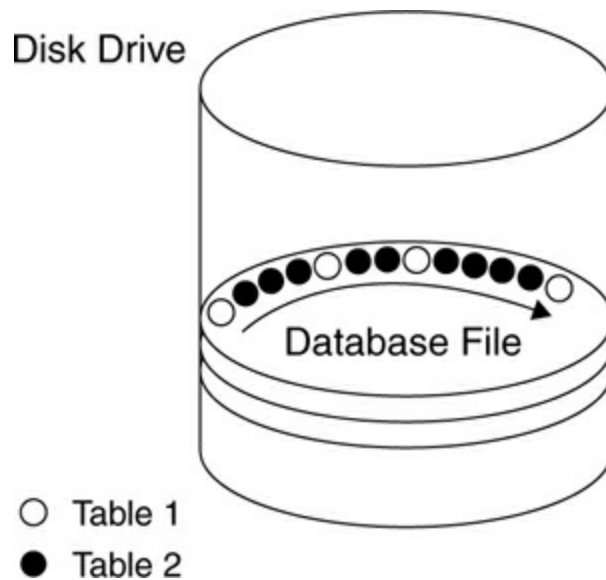


Figure 4.3. Interleaving table data

Different DBMS products support interleaving in different ways. Oracle uses a cluster to support a form of interleaving. For other DBMSs, you may need to develop scripts to organize, sort, and interleave the data before loading it. At any rate, interleaving is useful in a small number of cases and only when the predominant access to the two tables is by means of particular join criteria.

Denormalization

Normalization is the process of putting each fact in the most appropriate place. A normalized database implementation minimizes integrity problems and optimizes updating, perhaps at the expense of retrieval. When a fact is stored in only one place, retrieving many different but related facts usually requires going to many different places. This can slow the retrieval process. Updating is quicker, however, because the fact you're updating exists in only one place.

Most applications require very rapid data retrieval. Some applications require specific tinkering to optimize performance at all costs. To accomplish this, sometimes the decision is made to denormalize the physical database implementation. Just as normalization is the process of assembling data in an organized manner to eliminate redundancies, denormalization is the process of deliberately introducing redundancy to your data. In other words, denormalization can be thought of as the process of putting one fact in numerous places. This can have the effect of speeding up the data retrieval process, usually at the expense of data modification.

Denormalization can be thought of as the process of putting one fact in numerous places.

When to Denormalize

Of course, you should never denormalize data unless a performance need arises or your knowledge of the way your DBMS operates overrides the benefits of a normalized implementation. Many DBMS products have specific deficiencies and inefficiencies that may necessitate denormalizing for performance reasons. Therefore, denormalization is not necessarily a bad decision if implemented wisely. You should always consider the following issues before denormalizing:

- Can the system achieve acceptable performance *without* denormalizing?
- Will the performance of the system *after* denormalizing still be unacceptable?
- Will the system be less reliable due to denormalization?

If the answer to any of these questions is yes, you should avoid denormalization because the benefits typically will not exceed the cost. If, after considering these issues, you decide to denormalize, be sure to adhere to

the general guidelines that follow.

If enough disk space is available, consider creating two sets of tables: one set fully normalized and another denormalized. Populate the denormalized versions by querying the data in the normalized tables and loading or inserting it into the denormalized tables. Your application can access the denormalized tables in a read-only fashion and achieve performance gains, while at the same time modifying the normalized version and avoiding integrity problems in the base data. However, it is important to set up a controlled and scheduled population function to synchronize the normalized table with the denormalized.

If sufficient disk space is not available for two complete sets of tables, implement only the denormalized tables and maintain them programmatically. Be sure to update each denormalized table representing the same entity at the same time, or use database triggers to keep the redundant data synchronized.

When a column is replicated in many different tables, always update it everywhere simultaneously—or as close to simultaneously as possible—given the physical constraints of your environment. Once again, triggers can be helpful to accomplish this. If the denormalized tables are ever out of sync with the normalized tables, be sure to inform end users that batch reports and online queries may not contain sound data; if at all possible, this should be avoided.

Finally, be sure to design the application so that it can easily be converted from using denormalized tables to using normalized tables.

Every denormalization decision should be documented, including the reason behind the decision and the exact changes made from the normalized logical data model. Such a record will help to ensure that future database changes are made with appropriate knowledge. Documentation will also make it clear that you didn't simply make a design or implementation mistake.

Every denormalization decision should be documented.

Remember that only one valid reason exists for denormalizing a relational design—to enhance performance. The following criteria can be used to help identify potential denormalization candidates:

- Numerous critical queries or reports require data from more than one table—in other words, joins are required. If these types of requests need to be processed in an online transaction environment, denormalization may be able to improve performance.
- Repeating groups need to be processed in a group instead of individually.
- Many calculations need to be applied to one or many columns before queries can be successfully answered. Storing derived or precalculated data can reduce I/O and CPU usage upon retrieval and therefore be more efficient.
- Tables need to be accessed in different ways by different users during the same time frame.
- Many large primary keys exist that are clumsy to query and consume a large amount of disk space when carried as foreign key columns in related tables.
- Certain columns are queried a large percentage of the time, causing very complex or inefficient SQL to be used.

Of course, these situations do not always require denormalization, but they can be used as broad indications of when denormalization might be considered—which raises the question “When should the DBA denormalize a database design?” Although you might think it would be easier to denormalize at the very beginning of the physical design process, this is usually not the case. DBAs and application designers often decide to denormalize prematurely—before they obtain any concrete evidence for its necessity. Even though it may be difficult to retrofit a partially completed system to work on denormalized structures, it is almost never a good idea to denormalize before you’re sure that a normalized design will not perform adequately. Of course, sometimes a DBA will have direct experience for the specific application, DBMS, and version—and in those cases it may be acceptable to denormalize the physical design immediately. However, such a case is actually quite rare.

Be aware that each new RDBMS release usually brings enhanced performance and improved access options that may reduce the need for denormalization. However, most of the popular RDBMS products will require denormalized data structures on occasion. There are many different

types of denormalized tables that can resolve the performance problems caused when accessing fully normalized data. The following sections detail the different types and advise on when to consider implementing them.

Never, under any circumstances, should you attempt to create a denormalized logical data model. The logical data model should always be completely normalized. The physical implementation of the database can differ from the data model, but the model should always be fully normalized, and above all, all physical variations from the data model should be documented.

Never attempt to create a denormalized logical data model.

Prejoined Tables

If two or more tables need to be joined on a regular basis by an application, but the cost of the join is prohibitive, consider creating *prejoined tables*. The prejoined tables should

- Contain no redundant columns
- Contain only those columns absolutely necessary to meet the processing needs of the application
- Be created periodically, using SQL to join the normalized tables

The benefit of prejoining is that the cost of the join will be incurred only once—when the prejoined tables are created. A prejoined table can be queried very efficiently because every new query does not incur the overhead of the join process.

However, the negative aspect of prejoining, as with most forms of denormalization, is the difficulty of keeping the data accurate. Prejoined tables may quickly get out of sync with the independent tables from which they were created. For this reason, prejoined tables are more useful for relatively static data than for more dynamic data.

Report Tables

Oftentimes it is impossible to develop an end user report using only SQL. Such a report may require special formatting or data manipulation. If certain critical or highly visible reports of this nature are required to be viewed in an online environment, consider creating a table that represents the report. This table can then be queried using stand-alone SQL in a query tool or reporting

facility. The data for the report should be created by the appropriate mechanism (application program, 4GL, SQL, etc.) in a batch environment and then loaded into the report table in sequence. The report table should

- Contain one column for every column of the report
- Have its rows physically sequenced in the order in which they should appear on the report so sorting is not required
- Not subvert relational tenets (such as 1NF and atomic data elements)

The report table should contain one column for every column of the report.

Report tables are ideal for carrying the results of multiple joins and outer joins, correlated subqueries, or other complex SQL statements. If a complex query is coded, run, and then loaded into a table, a simple SELECT statement can be used to retrieve the results, instead of the complex (and perhaps slower) query that was used to populate the report table.

Mirror Tables

If an application system is very active, it may be necessary to split processing into two (or more) distinct components. Such a split will result in the creation of duplicate, or mirror, tables. For example, consider an application system with very heavy online traffic during the morning and early afternoon hours. This traffic consists of both queries and data modifications. Decision support processing is also performed on the same application tables during the afternoon. The production work in the afternoon always seems to disrupt the decision support processing, causing frequent time-outs and deadlocks.

These disruptions could be corrected by creating mirror tables—a foreground set of tables for the production traffic and a background set of tables for decision support. A mechanism to periodically migrate the foreground data to background tables must be established to keep the application data synchronized. One such mechanism could be a batch job executing UNLOAD and LOAD utilities. Another possibility is to use built-in replication and propagation software, if the DBMS supports such functionality. At any rate, the data synchronization should be done as often as necessary to sustain the effectiveness of the decision support processing.

It is important to note that the access needs of decision support are often considerably different from the access needs of the production environment.

Therefore, different decisions about data definition, such as indexing and clustering, may be made for the mirror tables.

In addition, simple mirror tables may not be sufficient for your decision support needs. Perhaps you will need to create a full-blown data warehouse environment. A data warehouse is just a relational database that is specifically designed or denormalized for decision support and analytical queries.

Split Tables

If separate pieces of one normalized table are accessed by different and distinct groups of users or applications, consider splitting the table into two (or more) denormalized tables—one for each distinct processing group. The original table can also be maintained if other applications access the entire table. In this scenario the split tables should be handled as a special case of mirror tables. If an additional table is not desired, a view joining the tables could be provided instead.

Tables can be split in one of two ways: vertically or horizontally. A *vertically split table* separates the columns of a table into separate tables: One set of columns is placed in a new table and the remaining columns are placed in another new table. The primary key columns are placed in both of the new tables. Designate one of the two new tables as the parent table for the purposes of referential integrity unless the original table still exists, in which case the original table should be the parent table in all referential constraints. If the original table still exists and the split tables are read only, don't set up referential integrity for the split tables. Because the split tables are derived from a referentially intact source, referential integrity is unnecessary.

Tables can be split in one of two ways: vertically or horizontally.

If you are splitting a table vertically, always include one row per primary key in each split table. Do not eliminate rows from any of the new tables for any reason. If rows are eliminated, the update process and any retrieval process that must access data from both tables will be unnecessarily complicated.

A *horizontally split table* separates the rows of a table into separate tables. To split a table horizontally, rows are classified into groups via key ranges.

The rows from one key range are placed in one table, those from another key range are placed in a different table, and so on. The columns of horizontally split tables are the same. For horizontal splits, avoid duplicating rows in the new tables. To accomplish this, use the primary key to perform the split and ensure that each key value is assigned to only one of the new tables.

Splitting Long Text Columns

A special case of a vertical split can be used to break up very large text columns. For example, consider a table that stores item descriptions of merchandise. The description can be 100 characters long, but most processes require only the first ten characters. For example, consider this CREATE TABLE statement:

[Click here to view code image](#)

```
CREATE TABLE ITEM
  (ItemNum      integer      not null,
   ItemSize     CHAR(1),
   ItemColor    CHAR(10),
   ItemDescr    CHAR(100)
  ;
```

In such a case, you can split the table into two tables by splitting the description into two columns. One new column, maintained in the old table, would house the first 10 bytes of the description. The second column would be created in a new table with the primary key and the last 90 bytes of the description. For example:

[Click here to view code image](#)

```
CREATE TABLE ITEM
  (ItemNum      INTEGER      NOT NULL,
   ItemSize     CHAR(1),
   ItemColor    CHAR(10),
   ItemDescr    CHAR(10)
  ;
CREATE TABLE ITEM_DESC
  (ItemNum      INTEGER      NOT NULL,
   ItemDesc     CHAR(90)
  ;
```

The value of this type of denormalization is better I/O: More rows can be stored on each physical page because each row of the main table is smaller.

Only those tasks that require all 100 bytes would need to access both tables. Of course, there are variations on this type of denormalization. You might choose to store only the first 10 bytes in the main table, but all 100 bytes in the description table if the other columns do not need to be accessed when the description is accessed. On the other hand, if the description is very large, you might want to break it up into multiple rows in the description table. For example, if the description can be up to 10,000 bytes long but most are under 1,000 bytes, you would not want to create a character column of 10,000 bytes (even if the DBMS allowed you to do so). Instead, you could create a table such as this:

The value of this type of denormalization is better I/O.

[Click here to view code image](#)

```
CREATE TABLE ITEM_DESC
  (ItemNum      INTEGER      NOT NULL,
   ItemCtr      INTEGER      NOT NULL,
   ItemDesc     CHAR(100)
  );
```

In this example, the primary key of the description table is now the combination of ItemNum and ItemCtr, where ItemCtr is a counter of the number of rows of description stored for the ItemNum. This design breaks the description up into 100-byte chunks. For the largest values, 100 rows would be required to store all 10,000 bytes, but for most descriptions, 10 or fewer rows would be required.

Combined Tables

If tables exist with a one-to-one relationship, consider combining them into a single table. Of course, if each participant in the one-to-one relationship has different relationships to other tables, you will need to take that into account when denormalizing. Sometimes even one-to-many relationships can be combined into a single table, but the data update process will be significantly complicated because of the increase in redundant data.

Consider combining tables with a one-to-one relationship into a single table.

For example, consider an application with two tables: DEPT (containing

department data) and EMP (containing employee data). You might choose to denormalize by combining the two tables into a large table named, for example, EMP_AND_DEPT. This new table would contain all of the columns of both tables except for the redundant foreign key. So, in addition to all the employee information, all the department information would also be contained on each employee row. This will result in many duplicate instances of the department data. Combined tables of this sort can be considered prejoined tables and treated accordingly.

Tables with one-to-one relationships should always be analyzed to determine whether combination is useful. Sometimes the consolidation of a one-to-one relationship is normalization, not denormalization.

Redundant Data

Sometimes one or more columns from one table are accessed almost every time data is queried in another table. In such cases, consider appending the columns to the queried table as redundant data. If the table carries these additional columns, joins can be eliminated and performance perhaps improved. This should be attempted only if the normal data access performs insufficiently.

Consider, once again, the DEPT and EMP tables. If most of the employee queries require the name of the employee's department, the department name column could be carried as redundant data in the EMP table. The column should not be removed from the DEPT table, though. Columns to consider storing redundantly should exhibit the following characteristics:

- Only a few columns are necessary to support the redundancy.
- The columns should be stable, needing infrequent updates.
- The columns should be used by either a large number of users or a few very important users.

Repeating Groups

The normalization process transforms repeating groups into distinct rows instead of separate columns of the same row. Even though the normalization of repeating groups optimizes data integrity and update performance, it usually results in higher disk usage and less efficient retrieval. This happens because there are more rows in the table and more rows need to be read in order to satisfy queries that access the repeating group.

Sometimes by denormalizing such groups back into distinct columns, you can achieve significant performance gains. Nevertheless, these gains come at the expense of flexibility. For example, consider an application that stores repeating group information in a table such as this:

[Click here to view code image](#)

```
CREATE TABLE CUST_BALANCE
  (CustNum          INTEGER          NOT NULL,
   BalancePeriod   INTEGER          NOT NULL,
   Balance         DECIMAL(15,2),
   constraint PKCB PRIMARY KEY (CustNum, BalancePeriod)
  ;
```

This table can store an infinite number of balances per customer, limited only by available storage and the storage limits of the DBMS. If the decision were made to string the repeating group, Balance, out into columns instead of rows, a limit would need to be set for the number of balances to be carried in each row. Here is an example of this table after denormalization:

[Click here to view code image](#)

```
CREATE TABLE CUST_BALANCE
  (CustNum          INTEGER          NOT NULL,
   Period1_Balance DECIMAL(15,2),
   Period2_Balance DECIMAL(15,2),
   Period3_Balance DECIMAL(15,2),
   Period4_Balance DECIMAL(15,2),
   Period5_Balance DECIMAL(15,2),
   Period6_Balance DECIMAL(15,2),
   constraint PKCB PRIMARY KEY (CustNum)
  ;
```

In this example, only six balances may be stored for any one customer. The designer could just as easily have chosen to store eight, 12, or any arbitrary number of balances. The number six is not important, but the concept that the number of values is limited is important. This reduces the flexibility of data storage and should be avoided unless performance needs dictate otherwise.

Using the first design, six rows would need to be retrieved to obtain six balances. Using the second design, all six balances can be retrieved by reading one row. Therefore, the performance of retrieval may be better using the denormalized design. Before deciding to implement repeating groups as columns instead of rows, be sure the following criteria are met:

- The data is rarely or never aggregated, averaged, or compared within the row.
- The data occurs in a statistically well-behaved pattern.
- The data has a stable number of occurrences.
- The data is usually accessed collectively.
- The data has a predictable pattern of insertion and deletion.

If any of these criteria are not met, certain types of data retrieval may be difficult to code, making the data less available. This should be avoided because, in general, data is denormalized only to make it *more* readily available.

Derivable Data

If the cost of deriving data using complicated formulas is prohibitive, think about physically storing the derived data in a column instead of calculating it. For example, consider employee data that is scattered across multiple tables. Perhaps the database contains three columns in several tables that store employee compensation data. These columns are Salary, Bonus, and Commission. Furthermore, assume that, more often than not, queries require total compensation to be reported, which is the sum of these three columns. It might make sense to include a column in the main EMP table called TotalCompensation that is the sum of Salary, Bonus, and Commission, thereby avoiding a multitable join and a calculation. Even though this example shows a simple addition, certain business calculations can be quite complex, requiring a lot of I/O and CPU processing to accomplish. The more complex the calculation and the more resources it requires, the better the performance gain you can achieve by physically storing it in the database instead of calculating the value every time it is required.

It can make sense to include a derived column in a table to avoid multitable joins and calculations.

However, when the underlying values that constitute the calculated value change, it is imperative that the stored derived data also be changed; otherwise inconsistent information will be stored. Such incorrect data will adversely impact the usability, effectiveness, and reliability of the database. To avoid such problems, consider storing derived data only when the following criteria are met:

- The source data used for the derivation calculation is relatively static.
- The cost of performing the derivation calculation is quite high.
- The usage pattern of the source tables is such that recalculation can be performed quickly when the source data changes.

Sometimes it is not possible to immediately update derived data elements when the columns upon which they rely change. Such situations can occur when the tables containing the derived elements are offline or being operated on by a utility. Whenever possible, time the update of the derived data so that it occurs immediately when the source table is made available again. Under no circumstances should outdated derived data be made available for reporting and inquiry purposes.

Hierarchies

A *hierarchy* is a structure that is easy to support using a relational database, but it can cause data retrieval difficulties unless the DBMS supports SQL extensions for traversing the hierarchy. For this reason, applications requiring hierarchies frequently contain denormalized tables to speed up data retrieval.

Applications requiring hierarchies frequently contain denormalized tables to speed up data retrieval.

Most of us have encountered at least one hierarchy in our data-processing careers. Two common hierarchical structures are bill-of-materials applications and departmental organization systems. A bill-of-materials application typically records information about parts assemblies in which one part is composed of other parts, which can then be a component of yet another part. A departmental organization system typically records the departmental structure of an organization, indicating which departments report to which other departments. A typical implementation of a hierarchy table for departmental organization would be

[Click here to view code image](#)

```
CREATE TABLE DEPT
  (ParentDeptNum    INTEGER          NOT NULL,
   DeptName         CHAR(25) ,
   SupervisorNum    INTEGER,
   ReportsToDeptNum INTEGER,
   constraint PKDN  PRIMARY KEY (DeptNum) ,
   constraint FKCB  FOREIGN KEY (ReportsToDeptNum)
```

```
REFERENCES
```

```
DEPT
```

```
ON DELETE RESTRICT
```

```
;
```

To support such a hierarchy, a one-to-many relationship is set up for a single table. In this example the ReportsToDeptNum column is the foreign key that refers to the DeptNum primary key. Each department reports to only one department, but a department can have more than one department reporting to it.

Such a table represents an accurately normalized entity for storing a hierarchy. The complete hierarchy can be rebuilt with the proper data retrieval instructions. However, consider the difficulty of writing SQL to query this table and report on the departmental organization. It is impossible to accomplish such a task using only SQL unless you have some guiding knowledge of the number of levels of reporting that may exist or your DBMS supports SQL extensions for traversing hierarchies.

A very effective way to denormalize a hierarchy is to create speed tables. The *speed table* contains a pretraversed hierarchy for easy retrieval. Such a speed table is shown here:

A very effective way to denormalize a hierarchy is to create speed tables.

[Click here to view code image](#)

```
CREATE TABLE DEPT
  (DeptNum          INTEGER          NOT NULL,
   ChildDeptNum    INTEGER          NOT NULL,
   Level           INTEGER,
   Detail          CHAR(1),
   DeptName        CHAR(25),
   SupervisorNum   INTEGER,
   constraint PKDN PRIMARY KEY (DeptNum, ChildDeptNum)
);
```

The speed table contains a row for every dependent ChildDeptNum, not just immediate dependents. The primary key for the speed table is the combination of DeptNum and ChildDeptNum. Two additional columns are provided:

- A column named Level contains a numeric value indicating the level within the hierarchy for the ChildDeptNum. For example, if the child resides two levels down in the hierarchy from the parent, Level will contain the value 2.
- A column named Detail contains “Y” if the ChildDeptNum is at the very bottom of the hierarchy, and “N” otherwise.

The speed table must be created programmatically—it cannot be generated using SQL.

Special Physical Implementation Needs

Sometimes the requirements of the database and the physical implementation details of the DBMS will not mix for good performance. For example, some DBMS products have limitations on the physical block sizes that can be specified for database files. At times, the row size of a table in a logical data model may require a very large block size because it will not fit completely in a smaller block size. Some DBMS products treat large block sizes inefficiently. In that case, you may want to denormalize the table by breaking apart the row so it will fit in a smaller block size. This is just one example of a physical DBMS implementation detail that may call for denormalization.

Denormalization Summary

We have discussed ten different types of denormalization. [Table 4.1](#) summarizes the types of denormalization that are available with a short description of when each type is useful.

Table 4.1. Types of Denormalization

Denormalization	Description
Prejoined tables	Used when the cost of joining is prohibitive
Report tables	Used when specialized critical reports are needed
Mirror tables	Used when tables are accessed concurrently by different types of environments
Split tables	Used when distinct groups use different parts of a table
Combined tables	Used to consolidate one-to-one or one-to-many relationships into a single table
Redundant data	Used to reduce the number of table joins required
Repeating groups	Used to reduce I/O and (possibly) storage usage
Derivable data	Used to eliminate calculations and algorithms
Speed tables	Used to make the processing of hierarchies more efficient
Physical denormalization	Used to optimize for specialized physical DBMS characteristics

The decision to denormalize should never be made lightly, because it can cause integrity problems and involve a lot of administration. Additional administration tasks include

- Documenting every denormalization decision
- Ensuring that all data remains valid and accurate
- Scheduling data migration and propagation jobs
- Keeping end users informed about the state of the tables
- Analyzing the database periodically to decide whether denormalization is still required

If a database has been denormalized, the data and environment should be regularly reviewed whenever hardware, software, and application requirements change. Any change can alter the need for denormalization. To verify whether denormalization is still a valid decision, ask the following questions:

Any change can alter the need for denormalization.

- Have the processing requirements changed for the application such that the join criteria, timing of reports, and/or transaction throughput requirements no longer dictate a denormalized database?
- Did a new release of the DBMS enhance performance? For example, did the introduction of a new join method or performance technique undo the need for prejoined tables?

- Did a new hardware release change performance considerations? For example, does the upgrade to a new, faster processor provide additional CPU resources so that denormalization is no longer necessary? Or did the addition of memory enable faster data access so that data can be physically normalized?

In general, periodically test whether the extra cost related to processing with normalized tables justifies the benefit of denormalization. You should measure the following criteria:

- I/O saved
- CPU saved
- Complexity of data modification
- Cost of returning to a normalized design

Always remember that denormalization is implemented to enhance performance. If the environment changes, it is only reasonable to reevaluate the denormalization decision. Also, it is possible that, given a changing hardware and software environment, denormalized tables may be causing performance degradation instead of performance gains. Simply stated, always monitor and periodically reevaluate all denormalized applications.

Denormalization is implemented to enhance performance.

Views

Another aspect of physical database design is the creation of *database views* to support specific application data requirements. Views are not required to access a physical database, but they can be helpful to support specific application and user requirements. You can think of a view as a way of turning a SELECT statement into a “table” that is accessible using SQL. Therefore, a view can be considered a logical table. No physical structure is required of a view; it is a representation of data that is stored in other tables (or other views). As shown in [Figure 4.4](#), the data “in the view” is not stored anywhere and physically exists only in the underlying tables. Views can also be based on other views.

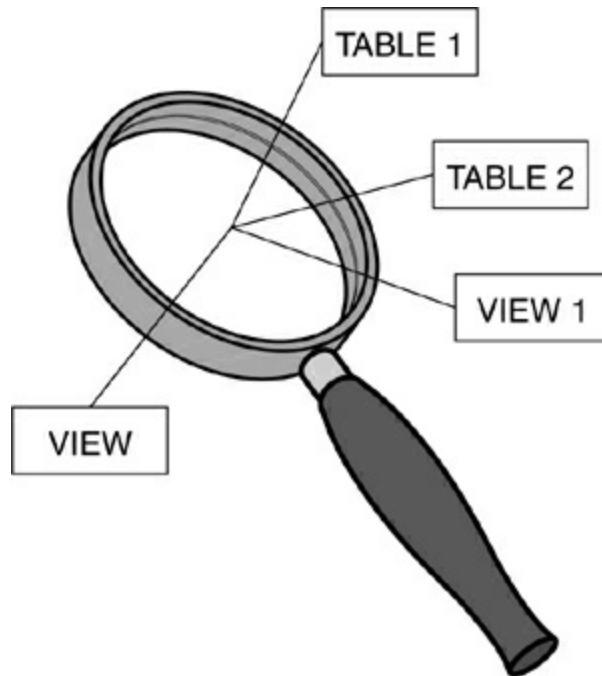


Figure 4.4. What is a view?

Views are flexible and can consist of any combination of the following:

- *Rows from tables.* These can be a subset of rows from a single table, all rows from a single table, a subset of rows from multiple tables, or all rows from multiple tables.
- *Rows from views.* These can be the same combinations as listed for tables.
- *Columns from tables.* These can be a subset of columns from a single table, all columns from a single table, a subset of columns from multiple tables, or all columns from multiple tables.
- *Columns from views.* These can be the same combinations as listed for tables.

Views should be created based on their usefulness to application development and ad hoc query users. There are six basic uses for which views excel. Views can allow you to

- Provide row- and column-level security
- Ensure efficient access paths
- Mask complexity from the user
- Ensure proper data derivation

- Rename tables
- Rename columns

At any rate, be sure to document the intended purpose for every view created so that future structural changes to database tables can be promulgated to any views accessing those changed tables.

Data Definition Language

All physical database objects are created using SQL Data Definition Language, or DDL for short. The basic components of DDL are the CREATE, ALTER, and DROP statements. Appropriately enough, CREATE is used to initially create a database object. Changes to the database object once it has been created can be made using the ALTER statement. But the ALTER statement cannot necessarily be used to change any and every aspect of a database object (which is covered in more detail in [Chapter 7, “Database Change Management”](#)). Finally, the DROP statement is used to remove a database object from the system.

All physical database objects are created using SQL Data Definition Language.

Many DBMS products provide a graphical interface for creating and changing database objects. If your DBMS provides such an interface, you may be able to create a physical database without learning the specifics of DDL syntax. I do not recommend this for DBAs, because sometimes the graphical interface does not support all of the syntax and options for every database object. An informed DBA is an effective DBA—and unless you verify that the graphical interface supports every DDL option, you will be better off learning and using DDL statements.

Temporal Data Support

The need to store and access noncurrent data is a common requirement for some types of applications. Many types of data change over time, and different users and applications have requirements to access the data at different points in time. For some, the current, up-to-date values for the data are sufficient. But for others, the ability to access earlier versions of the data is needed.

In traditional relational database systems, various approaches have been used to store and access temporal data. Separate history tables are one approach (perhaps augmented with triggers); snapshot tables are another; yet an additional approach is to custom-build time sensitivity into the tables and queries. None of these approaches is ideal, especially with the fast pace of business and the escalating need to comply with legal and regulatory requirements. Many applications need to provide real-time access to noncurrent data.

A traditional database stores data implied to be valid at the current point in time; it does not track the past or future states of the data. But some database systems (for example, DB2 for z/OS) support built-in temporal features. Temporal support makes it possible to store different database states and to query the data as of those different states. This is accomplished by attaching a time period to the data to indicate when it was valid or changed in the database.

A Temporal Example

Consider an insurance company that sells policies to its customers. The terms of any specific insurance policy are valid over a period of time. After that period of time, customers can choose to decline further coverage, continue with the existing coverage, or modify the terms of their coverage. So at any specific point in time, the terms of the customers' policies can differ.

Over time, customers make claims against their policies. This claim information needs to be stored, managed, and analyzed. Accident histories for customers are also important pieces of data with a temporal element.

Now consider the complexity inherent in trying to develop not only a database design that accommodates changing policies, claims, and historical details but also enables queries such that a user might access a customer's coverage at a given point in time. In other words, what policies were in effect for that customer as of, say, April 15, 2012? Or on any other date during which the customer had coverage?

This concept of business time can become quite complex. Consider the situation in which a customer has multiple policies that expire on different schedules. Add the possibility for periods in which coverage lapses. And the database does not remain static; it grows in size, queries become more complex, and so on.

Insurance is but one example. Many other types of applications exist for which temporal support would be useful, for example, financial applications, credit history, personnel management, transportation applications, reservation systems, and medical information management, to name a few.

Business Time and System Time

There is another concept of temporal data that might need to be factored in as well. Instead of business time, you also might need to track and manage system time. For example, a regulatory mandate might require you to track any changes to a particular piece of data. This is common with personally identifiable information (PII), such as a Social Security number or phone number. Support for managing system changes enables users to query the database as of a point in time, returning the value of the data as of that time frame.

The business time indicates the period during which the data is accurate with respect to the world. The system time indicates the period during which the data is stored in the database. These two time periods do not need to be the same for a single piece of data. Suppose, for example, that you want to store temporal information about twentieth-century events. The valid business time for this data would be within the range of 1900 through 1999. But if you were to add the information to the database now, perhaps the valid system time for the data would be at some point in 2013. So, of course, you might need to support both business temporal data and system temporal data in the same table. This is called *bitemporal support*.

So, to define terms a bit more precisely: *Business time*, also referred to as valid time or application time, specifies when the facts stored in the database are true for the real world. These are the dates of interest to the business user interacting with the data. Business time should be associated only with data that has a business need for temporal tracking.

System time, also referred to as transaction time, denotes the time when the fact became current in the database. System time can be used to track the insertion and modification history of the data. Unlike business time, transaction time can be associated with any database entity. It might be useful, depending upon regulatory or industry requirements, to track when any specific piece of data changes. Of course, you do not want to impose a system time on all your data just because you can due to the overhead of

doing so.

A bitemporal table is defined with both a system period and a business period. With a bitemporal table you can keep business (or application) information and system-based historical information.

Impact on DBA

DBAs must understand the temporal requirements of the database before it is implemented. The DBA should properly implement the business and system time using the built-in facilities of the DBMS. When this is done, temporal queries are easy to code using the AS OF syntax of SQL. For example, consider our discussion of tracking the insurance policies in effect for a customer as of April 15, 2012. For a temporal database, this can be simply using something like the following SQL:

[Click here to view code image](#)

```
SELECT  CustName, PolicyNo, BenefitSummary
FROM    InsurancePolicy
        FOR BUSINESS_TIME AS OF '2012-04-15'
WHERE   CustNo = ?
;
```

To modify temporal data you use standard INSERT, UPDATE, and DELETE statements, but you must understand the temporal aspect of the changes being made. For INSERT statements you must manipulate the date range correctly. For DELETE statements you must assure that the correct ranges are being removed. And UPDATE statements can be augmented with temporal clauses so that specific ranges are modified.

Of course, if the DBMS does not offer temporal support, the DBA and application team will need to build the tables with effective dates and ensure that the application code is created to modify the time spans appropriately. Additionally, queries against the tables become more complex, requiring BETWEEN or greater-than/less-than predicates for the date ranges.

This discussion of temporal data is necessarily brief. For a more detailed account consult Johnston and Weis, *Managing Time in Relational Databases*.

Summary

A logical data model should be used as the blueprint for designing and creating a physical database. But the physical database cannot be created with

a simple logical-to-physical mapping. Many physical design decisions need to be made by the DBA before implementing physical database structures. Many times this will necessitate deviating from the logical data model. But such deviation must be based on in-depth knowledge of the DBMS and the physical environment in which the database will exist.

Review

1. Describe the first, simple steps to convert a logical data model to a physical database.
2. What is the only reason to denormalize a physical data model?
3. Under what circumstances should a bitmap index be considered instead of a b-tree index?
4. Which types of data access will benefit from data clustering?
5. Cite five reasons for creating a database view.
6. A referential constraint is created by specifying a _____ key in the dependent table that refers back to the primary key in the parent table.
7. Describe how a relational database uses indexes.
8. Why might the order in which columns are created in a table be important for physical database design?
9. When should you consider physically storing derived data in a database?
10. If indexes are beneficial to performance, why not create every possible index conceivable just to be on the safe side?

Bonus Question

Review the very small data model shown in [Figure 4.5](#) and create a physical database implementation for the DBMS of your choice. Assume that there are approximately 25,000 students and that each student enrolls in three to five courses per semester. The most common query requirement is that students be able to create their course schedules. Be sure to specify every physical design decision and create sample DDL statements to implement a basic physical design. Further, indicate where additional information is required to make a physical design decision and why.

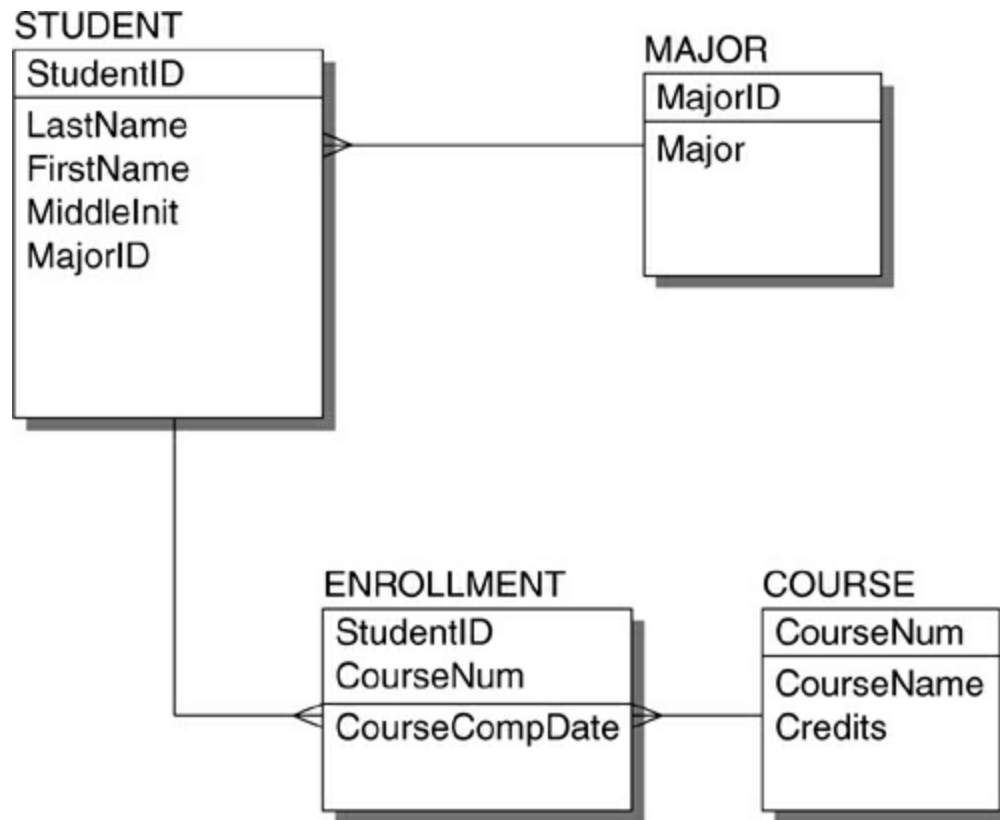


Figure 4.5. Logical data model

Suggested Reading

- Fleming, Candace, and Barbara von Halle. *Handbook of Relational Database Design*. Reading, MA: Addison-Wesley (1989). ISBN 0-201-11434-8.
- Harrington, Jan L. *Relational Database Design: Clearly Explained*. 2nd ed. San Francisco, CA: Morgan Kaufmann (2002). ISBN 1-55860-820-6.
- Hernandez, Michael J. *Database Design for Mere Mortals*. 2nd ed. Boston, MA: Addison-Wesley (2003). ISBN 0-201-75284-0.
- Hogan, Rex. *A Practical Guide to Database Design*. Englewood Cliffs, NJ: Prentice Hall (1990). ISBN 0-13-690967-1.
- Johnston, Tom, and Randall Weis. *Managing Time in Relational Databases*. Burlington, MA: Morgan Kaufmann (2010). ISBN 978-0-12-375041-9.
- Lahdenmaki, Tapio, and Michael Leach. *Relational Database Index Design and the Optimizers*. Hoboken, NJ: John Wiley & Sons (2005). ISBN 0-471-71999-4.

- Lightstone, Sam, et al. *Physical Database Design*. San Francisco, CA: Morgan Kaufmann (2007). ISBN 978-0-12-369389-1.
- Pascal, Fabian. *Practical Issues in Database Management*. Boston, MA: Addison-Wesley (2000). ISBN 0-201-48555-9.
- Perkinson, Richard C. *Data Analysis: The Key to Database Design*. Wellesley, MA: QED Information Sciences (1984). ISBN 0-89435-105-2.
- Riordan, Rebecca M. *Designing Effective Database Systems*. Boston, MA: Addison-Wesley (2003). ISBN 0-321-29093-3.
- Rishe, Naphtali. *Database Design*. New York: McGraw-Hill (1992). ISBN 0-07-052955-8.
- Stephens, Ryan K., and Ronald R. Plew. *Database Design*. Indianapolis, IN: SAMS Publishing (2001). ISBN 0-672-31758-3.
- Teory, Toby, et al. *Database Design: Know It All*. San Francisco, CA: Morgan Kaufmann (2002). ISBN 978-0-12-374630-6.

5. Application Design

Application design is more than just writing efficient database requests in application programs. Every aspect of the way the program is coded affects the usability and effectiveness of the application. Of course, application design includes database concerns such as interfacing SQL with traditional programming languages and the type of SQL to use. However, each application program must be designed to ensure the integrity of the data it modifies. Additionally, performance has to be treated as a design issue.

Performance has to be treated as a design issue.

At the forefront the DBA must promote the concept of application design based on thorough knowledge of the database. It is unacceptable to allow programmers to design and code applications without considering how the programs will perform as they interact with databases. Some uninformed organizations approach database application development with no proactive performance engineering: The assumption is that any performance problems can be resolved after development by the DBA. However, it may be impossible to tune an improperly designed application program without rewriting it. So why not write it correctly the first time?

The intent of this chapter is not to teach software development methodology or to provide an in-depth treatise on programming—nor is it a primer on SQL. The focus of discussion will be high-level application design issues that need to be understood when writing applications that use a database for persistent storage of data. All DBAs should understand the concepts in this chapter and be able to effectively communicate them to the developers in their organization.

[Chapter 12, “Application Performance,”](#) provides additional coverage of application performance issues as they pertain to database development.

Database Application Development and SQL

Designing a proper database application system is a complex and time-consuming task. The choices made during application design will impact the usefulness of the final, delivered application. Indeed, an improperly designed

and coded application may need to be redesigned and recoded from scratch if it is inefficient, ineffective, or not easy to use.

Designing a proper database application system is a complex task.

To properly design an application that relies on databases for persistent data storage, the system designer at a minimum will need to understand the following issues:

- How data is stored in a relational database
- How to code SQL statements to access and modify data in the database
- How SQL differs from traditional programming languages
- How to embed SQL statements into a host programming language
- How to optimize database access by changing SQL and indexes
- Programming methods to avoid potential database processing problems

In general, the developer must match the application development languages and tools to the physical database design and functionality of the DBMS. The first task is to master the intricacies of SQL.

SQL

Structured Query Language, better known as SQL (and pronounced “sequel” or “ess-cue-el”), is the de facto standard for accessing relational databases. All RDBMS products, and even some nonrelational DBMS products, use SQL to manipulate data.

SQL is the de facto standard for accessing relational databases.

Why is SQL so pervasive within the realm of relational data access? There are many reasons for SQL’s success. Foremost is that SQL is a high-level language that provides a greater degree of abstraction than traditional procedural languages. Third-generation languages, such as COBOL and C, and even fourth-generation languages, usually require the programmer to navigate data structures. Program logic must be coded to proceed record by record through data stores in an order determined by the application programmer or systems analyst. This information is encoded in the high-level language and is difficult to change after it has been programmed.

SQL, by contrast, is designed such that programmers specify *what* data is needed. It does not—indeed it cannot—specify *how* to retrieve it. SQL is coded without embedded data-navigational instructions. The DBMS analyzes each SQL statement and formulates data-navigational instructions “behind the scenes.” These data-navigational instructions are commonly called [access paths](#). A heavy burden is removed from the programmer by empowering the DBMS to determine the optimal access paths to the data. Because the DBMS better understands the state of the data it stores, it can produce a more efficient and dynamic access path to the data. The result is that SQL, used properly, provides a quicker application development and prototyping environment than is available with corresponding high-level languages. Furthermore, as the data characteristics and access patterns change, the DBMS can change access paths for SQL queries without requiring the actual SQL to be changed in any way.

SQL specifies what data is needed . . . not how to retrieve it.

Inarguably, though, the single most important feature that has solidified SQL’s success is its capability to retrieve data easily using English-like syntax. It is much easier to understand a query such as

```
SELECT    deptnum, deptname
FROM      dept
WHERE     supervisorum = '903';
```

than it is to understand pages and pages of C or BASIC source code, let alone the archaic instructions of Assembler. Because SQL programming instructions are easier to understand, they are easier to learn and maintain—affording users and programmers more productivity in a shorter period of time. However, do not underestimate SQL: Mastering all of its intricacies is not easy and will require much study and practice.

SQL also uses a free-form structure that makes it very flexible. The SQL programmer has the ability to develop SQL statements in a way best suited to the given user. Each SQL request is parsed by the DBMS before execution to check for proper syntax and to optimize the request. Therefore, SQL statements do not need to start in any given column and can be strung together on one line or broken apart on several lines. For example, the following SQL statement:

[Click here to view code image](#)

```
SELECT deptnum, deptname FROM dept WHERE supervisornum
= '903';
```

is exactly equivalent to the previous SQL statement shown. Another example of SQL’s flexibility is that the programmer can formulate a single request in a number of different and functionally equivalent ways—a feature that also can be very confusing for SQL novices. Furthermore, the flexibility of SQL is not always desirable, because different but logically equivalent SQL formulations can result in differing performance results. Refer to the sidebar “[Joins versus Subqueries](#)” for an example.

Finally, one of the greatest benefits derived from using SQL is its ability to operate on sets of data with a single line of code. Multiple rows can be retrieved, modified, or removed in one fell swoop by using a single SQL statement. This feature provides the SQL developer with great power but also limits the overall functionality of SQL. Without the ability to loop or step through multiple rows one at a time, certain tasks are impossible to accomplish using only SQL. Of course, as more and more functionality is added to SQL, the number of tasks that can be coded using SQL alone is increasing. For example, SQL can be used to create a stored procedure to perform many programming tasks that formerly required a traditional programming language to accomplish. Furthermore, most of the popular relational DBMS products support extended versions of SQL with procedural capabilities. [Table 5.1](#) details the most popular procedural SQL dialects.

Table 5.1. SQL Usage Considerations

DBMS	Procedural SQL Dialect
Oracle	PL/SQL
Microsoft SQL Server	Transact-SQL
Sybase Adaptive Server Enterprise	Transact-SQL
DB2	SQL Procedure Language

Joins versus Subqueries

One example of SQL’s flexibility is the way in which a single statement can access data from multiple tables. SQL provides two methods: [joins](#) and [subqueries](#). However, a subquery can be converted to an equivalent join.

The concept behind both types of queries is to retrieve data from multiple tables based on search criteria matching data in the tables.

Consider the following two SQL statements. The first one is a subquery, where a query is embedded within another query. The second query is a join, where two tables are specified in the FROM clause of a single SELECT statement.

[Click here to view code image](#)

```
SELECT empno, firstname, lastname
FROM employee
WHERE workdept IN
      (SELECT deptno
       FROM department
       WHERE deptname = 'DBA');

SELECT empno, firstname, lastname
FROM employee,
      department
WHERE workdept = deptno
AND deptname = 'DBA';
```

Both queries return information about employees who work in the database administration department. The results returned by both queries will be the same, but the performance may vary significantly depending on the DBMS in use, the indexes that are defined for each table, and the complexity of the query itself.

Set-at-a-Time Processing and Relational Closure

Every operation performed on a relational database operates on a table (or set of tables) and results in another table. This feature of relational databases is called *relational closure*. All SQL data manipulation operations—that is, SELECT, INSERT, UPDATE, and DELETE statements—are performed at a set level. One retrieval statement can return multiple rows; one modification statement can modify multiple rows.

All SQL data manipulation operations are performed at a set level.

To clarify the concept of relational closure, refer to [Figure 5.1](#). A database user initiates SQL requests. Each SQL statement can access one or many

tables in the database. The SQL statement is sent to the DBMS, whereupon the query is analyzed, optimized, and executed. The DBMS formulates the access path to the data, and upon completion of the request the desired information is presented to the user as a set of columns and rows—in other words, a table. The result will consist of one or more columns with zero, one, or many rows. Because SQL performs set-level processing, the DBMS operates on a set of data, and a set of data is always returned as the result. Of course, the results set can be empty, or it can contain only one row or column. The relational model and set-level processing are based on the mathematical laws of *set theory*, which permit empty and single-valued sets.

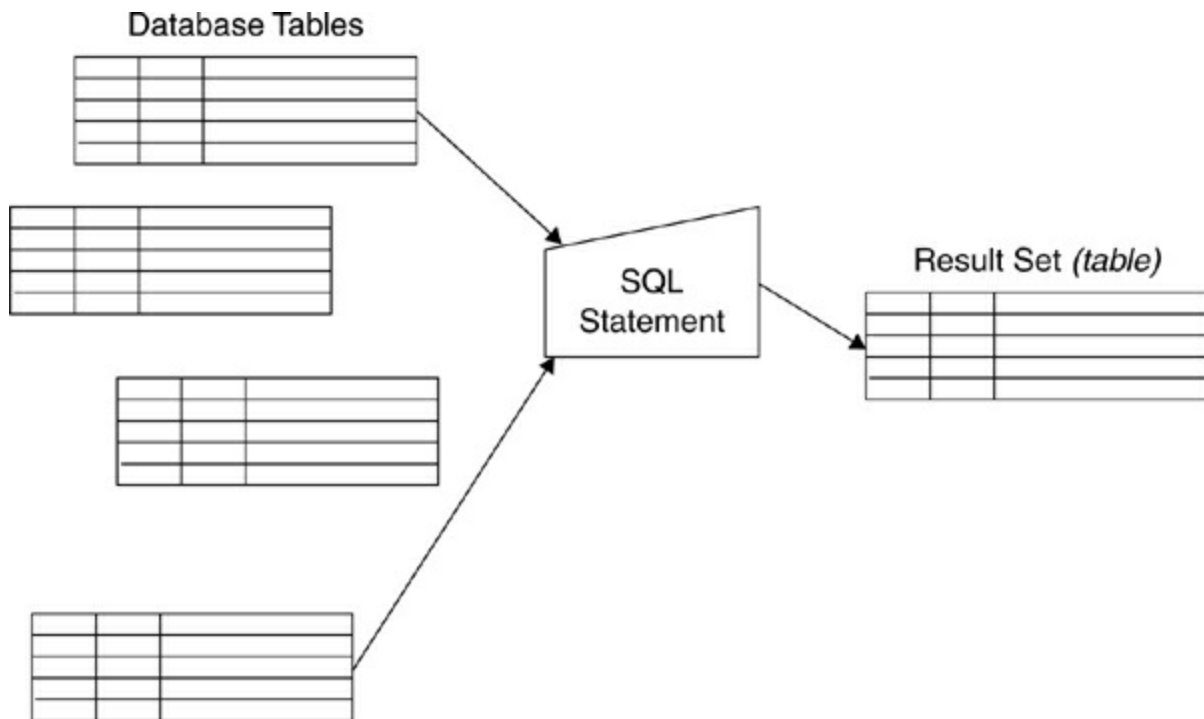


Figure 5.1. Relational closure

Application developers face a potential problem when using relational databases because of the set-at-a-time nature of SQL. Most programming languages operate on data one record at a time. When a program requires relational data, though, it must request the data using SQL. This creates an impedance mismatch. The program expects data to be returned a single row at a time, but SQL returns data a set at a time. There are different ways to get around this mismatch, depending on the DBMS, programming language, and environment. Most DBMS products provide a feature called a [*cursor*](#) that accepts the input from a SQL request and provides a mechanism to fetch

individual rows of the results set. Some programming environments and fourth-generation languages automatically transform multirow sets to single rows when communicating with the DBMS.

Furthermore, most programmers are accustomed to hard-wiring data-navigational instructions into their programs. SQL specifies what to retrieve but not how to retrieve it. The DBMS determines how best to retrieve the data. Programmers unaccustomed to database processing are unlikely to grasp this concept without some training. At any rate, programmers will need to be trained in these high-level differences between non-database and database programming techniques. This job usually falls upon the DBA or other highly skilled database technician within the organization. Of course, there are many more details at a lower level that the database programmer needs to know, such as SQL syntax, debugging and testing methods, optimization techniques, and program preparation procedures (compilation, bind, etc.).

The DBMS determines how best to retrieve the data.

Embedding SQL in a Program

Most database applications require a host programming language to use SQL to communicate with the database. A wide range of programming languages can be used with SQL, from traditional languages such as COBOL, FORTRAN, and Assembler to more modern languages such as C/C++, Java, PHP, and Visual Basic. Your choice of host programming language can impact the way you will have to code SQL. For example, SQL is embedded directly into a COBOL program, whereas a language like C requires an API such as ODBC to issue SQL statements.

The choice of development language will likely be limited to just a few at your shop. You should attempt to minimize the number of different languages you use, because it will make supporting and maintaining your applications easier. Furthermore, such limitations will make it easier for DBAs to administer and optimize the database environment. A DBA should be capable of reading and understanding program code for each language used to access databases within the organization.

Minimize the number of different languages you use.

Some application development projects use an *IDE (integrated development environment)* or *code generator* to create programs from program specifications. Exercise caution when using this approach: Don't allow the code generator to create SQL for you without first testing it for efficiency. Poor performance can result when using a code generation tool because these tools often have very little knowledge of the DBMS you are using. In most cases the code generator is designed to work for multiple DBMS products and is therefore not optimized for any of them. Test the SQL that is generated and, if necessary, modify the generated SQL or build indexes to optimize the generated SQL before moving the programs to a production environment. The DBA should be responsible for implementing a procedure to ensure that such SQL performance testing occurs.

SQL Middleware and APIs

Application programs require an *interface* for issuing SQL to access or modify data. The interface is used to embed SQL statements in a host programming language, such as COBOL, Java, C, or Visual Basic. Standard interfaces enable application programs to access databases using SQL. There are several popular standard interfaces or APIs for database programming, including ODBC, JDBC, SQLJ, and OLE DB.

One of the most popular SQL APIs is Open Database Connectivity (ODBC). Instead of directly embedding SQL in the program, ODBC uses *callable routines*. ODBC provides routines to allocate and deallocate resources, control connections to the database, execute SQL statements, obtain diagnostic information, control transaction termination, and obtain information about the implementation. ODBC is basically a call-level interface (CLI) for interacting with databases. The CLI issues SQL statements against the database by using procedure calls instead of direct embedded SQL statements.

ODBC is basically a call-level interface for interacting with databases.

Microsoft invented the ODBC interface to enable relational database access for Microsoft Windows programming. However, ODBC has become an industry-standard CLI for SQL programming. Indeed, every major DBMS today supports ODBC.

ODBC relies on *drivers*, which are optimized ODBC interfaces for a particular DBMS implementation. Programs can make use of the ODBC drivers to communicate with any ODBC-compliant database. The ODBC drivers enable a standard set of SQL statements in any Windows application to be translated into commands recognized by a remote SQL-compliant database.

Another popular SQL API is Java Database Connectivity (JDBC). JDBC enables Java to access relational databases. Similar to ODBC, JDBC consists of a set of classes and interfaces that can be used to access relational data. There are several types of JDBC middleware, including the JDBC-to-ODBC bridge, as well as direct JDBC connectivity to the relational database. Anyone familiar with application programming and ODBC (or any call-level interface) can get up and running with JDBC quickly. (Refer to [Chapter 21](#), “[Database Connectivity](#),” for a discussion of the various types of JDBC drivers.)

Another way to access databases from a Java program is by using SQLJ. SQLJ enables developers to directly embed SQL statements in Java programs, thus providing static SQL support to Java. A precompiler translates the embedded SQL into Java code. The Java program is then compiled into bytecodes, and a database bind operation creates packaged access routines for the SQL.

SQLJ enables developers to embed SQL statements in Java programs.

OLE DB, which stands for Object Linking and Embedding Database, is an interface that is based on the COM architecture. It is a low-level interface to data. OLE DB provides applications with uniform access to data stored in diverse information sources. It allows greater flexibility than ODBC because it can be used to access both relational and nonrelational data. OLE DB presents an object-oriented interface for generic data access. OLE DB is conceptually divided into *consumers* and *providers*. The consumers are the applications that need access to the data, and the providers are the software components that implement the interface and thereby provide the data to the consumer.

SQL Server 2012¹ is planned to be the final version to include an OLE DB provider for SQL Server. However, Microsoft has indicated that support will

continue for at least seven years.

COM is Microsoft's component-based development architecture. Using COM, developers can create application components that can be pieced together to create application systems. The components can be written by different developers and need not be written using the same programming language. ADO (which stands for ActiveX Data Objects) is a set of software components that programmers can use to access data and data services.

Both COM and ADO predate the .NET framework but have been adapted for use by .NET.

Application Infrastructure

Application infrastructure is the combined hardware and software environment that supports and enables the application. The application infrastructure will vary from organization to organization, and even from application to application within an organization. The application infrastructure provides the foundation for building, deploying, and managing applications with high performance, security, and control.

From a hardware perspective, the application infrastructure includes the servers, clients, and networking components. From a software perspective, things are a bit more difficult to nail down. Software components of an application infrastructure can include database servers, application servers, Web servers, transaction managers, and development frameworks. Some of the key functionality of application infrastructure includes transaction management, clustering, reliable application-to-application messaging, system management, advanced application development tools, proprietary access, and interoperability with legacy technologies.

From a mainframe perspective, the application infrastructure may consist of IBM z Series hardware running z/OS, DB2, CICS, with application programs written in COBOL. Typically, applications consist of both batch and online workload. A modern mainframe infrastructure adds interfaces to non-mainframe clients, as well as WebSphere Application Server and Java programs. Most new mainframe development uses IDEs to code modern applications instead of relying upon COBOL programmers.

Most modern, distributed, non-mainframe application development projects typically rely upon application development frameworks. The two most commonly used frameworks are Microsoft .NET and J2EE.

.NET

The Microsoft .NET framework provides a comprehensive development platform for the construction, deployment, and management of applications. The .NET framework provides CLR (common language run time) and a class library for building components using a common foundation. This offers benefits to developers such as support for standard practices, extensibility, and a tightly integrated set of development tools.

The .NET framework consists of multiple major components in addition to the CLR and class library. From a data perspective, the most important component is ADO.NET, which provides access to data sources such as a database management system.

See [Figure 5.2](#) for a depiction of the .NET framework and its components.

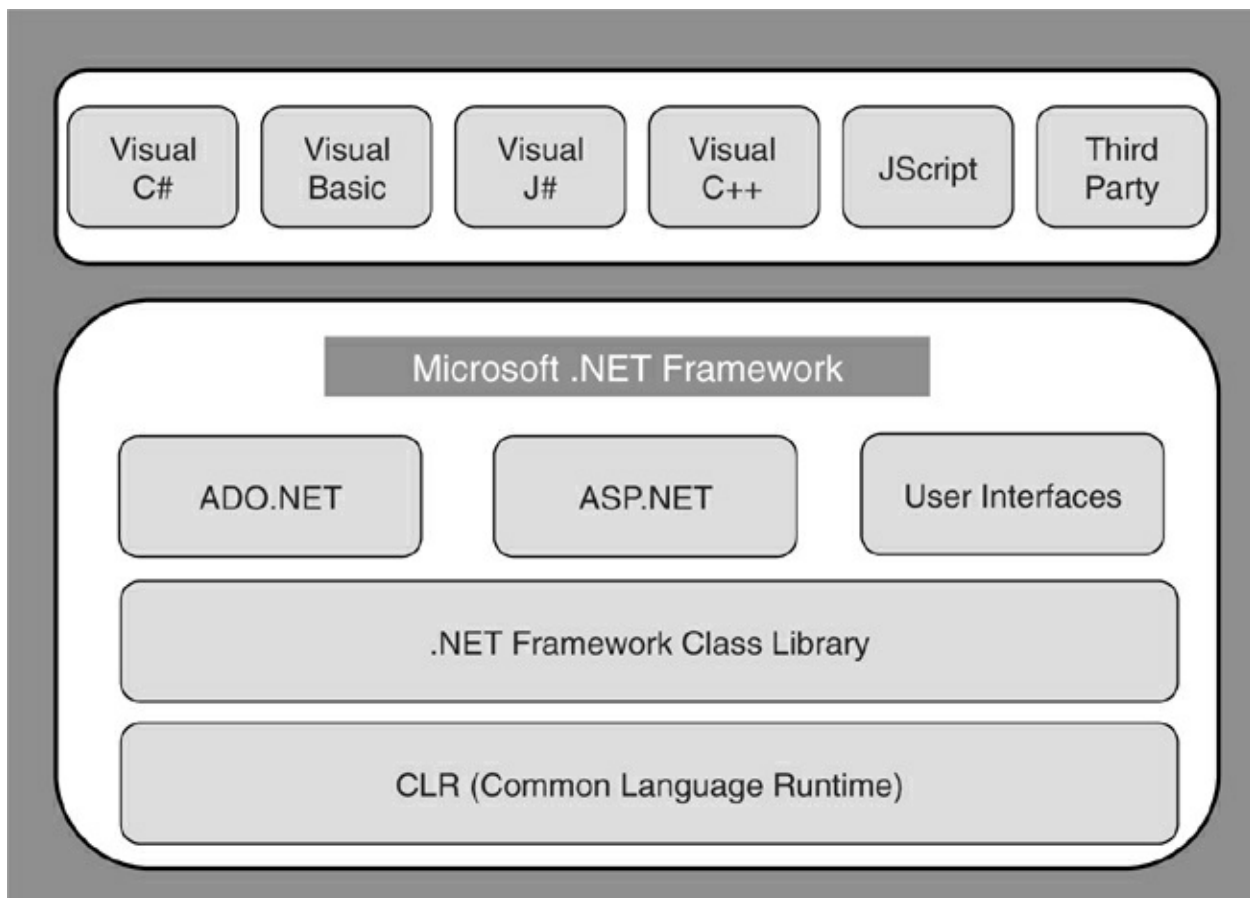


Figure 5.2. The .NET framework

ADO.NET is composed of a series of technologies that enable .NET developers to interact with data in standard, structured, and predominantly

disconnected ways. Applications that use ADO.NET depend on .NET class libraries provided in DLL files. ADO.NET manages both internal data (created in memory and used by the program) and external data (in the database). It provides interoperability and maintainability through its use and support of XML, simplified programmability with a programming model that uses strongly typed data, and enhanced performance and scalability.

J2EE and Java

The Java 2 Platform, Enterprise Edition (J2EE), is a set of coordinated specifications and practices that together enable solutions for developing, deploying, and managing multitier enterprise applications. The J2EE platform simplifies enterprise applications by basing them on standardized, modular components. J2EE provides a complete set of services to those components and handles many details of application construction without requiring complex programming.

So J2EE is not exactly a software framework, but a set of specifications, each of which dictates how various J2EE functions must operate. Software conforming to the J2EE platform offers advantages such as “Write Once, Run Anywhere” portability, JDBC API for database access, CORBA technology for interaction with existing enterprise resources, and a security model for data protection. Building on this base, the Java 2 Platform, Enterprise Edition, adds full support for Enterprise JavaBeans components, Java Servlets API, JavaServer Pages, and XML technology.

See [Figure 5.3](#) for a depiction of a sample J2EE implementation. Additional information and clarification on J2EE can be found online at http://java.sun.com/j2ee/reference/whitepapers/j2ee_guide.pdf.

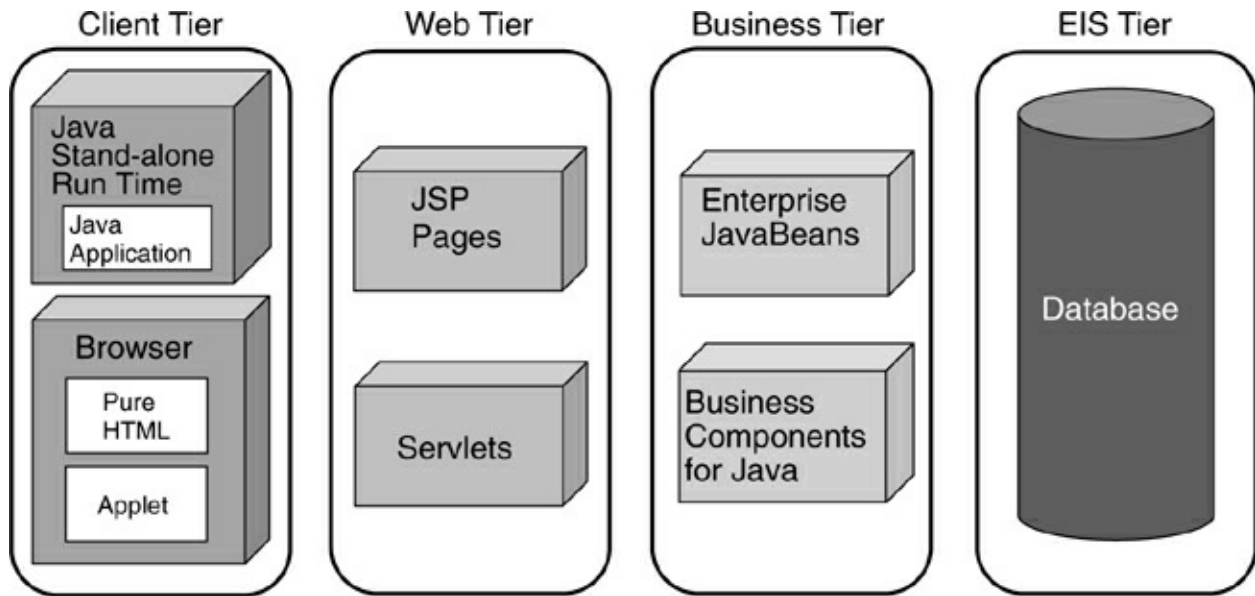


Figure 5.3. A sample J2EE implementation

Note that there is much more to Java than is covered in this section. Refer to the sidebar “[Java Program Types](#)” for a brief overview of the different types of Java programs that can be coded.

Java Program Types

There are three types of programs that you can implement when accessing data using Java: Java applets, Java servlets, and Java applications.

A *Java applet* is a small application program that must be downloaded before it can be run by a Java-enabled Web browser. Java applets reside on a Web server. When the Web server returns an HTML page that points to a Java applet, the Java-enabled Web browser requests the applet to be downloaded. After the applet is received at the browser, either the browser starts the applet internally or an external JVM executes it.

Applets typically perform simple operations, such as editing input data or controlling screen interaction, and provide other client functionality. Of course, Java applets can be written to perform more complex functionality, but to load and run non-Java code in the client requires *signed applets*, which have the authority needed to run code in the client machine.

Performance problems can arise because Java applets must be

downloaded before they can be run. The time to download the applet must be factored into its service levels. In general, Java applets are small, so the performance impact should be negligible. Nevertheless, even small downloads can be slow if there are network problems. Java applets can be cached by the Web browser, which diminishes the performance impact.

A *Java servlet* is basically server-side Java. A Java servlet runs on the Web server, just as an applet runs in the Web browser. Java servlets can be used to extend the functionality of the Web server. The Web server hands requests to the servlet, and the servlet replies to them. Servlets can be used instead of CGI applications.

Java servlets have security advantages over client-side Java applets. A servlet that runs on a Web server inside a firewall can control access to sensitive data and business logic. Java applets do not inherently provide these security capabilities.

A *Java application* program is basically the same as a program written in any other programming language. It can perform all of the tasks normally associated with programs, including many tasks that Java applets cannot perform. Furthermore, a Java application does not need a browser to be executed. It can be executed in a client or server machine.

Before choosing which Java development style to use, you must know the basics of the environment in which the program will be run. Ask the following questions when deciding what type of Java program is required for your development needs:

- How will the program be executed? Must it run over the Internet, as an intranet or extranet application, or merely as a stand-alone application?
- What is the business logic that this program must perform?
- How complicated is the program?
- How large (or small) is the program, and can it be quickly downloaded?
- What are the security requirements?
- Who are the target users and at what speed will they be connected to the Web?

Java applications, Java applets, and Java servlets are similar in nature, but a different method is used to invoke each of them. Java applets and

servlets are started from an HTML page. Java applications do not require a Web component but can be used as part of an intranet solution.

.NET versus J2EE

There is an ongoing debate as to the relative merits of .NET versus the J2EE platform. In reality, it is not an “either/or” decision that organizations are making, but a “both/and” decision. Although both are development platforms, the two are not interchangeable.

At a very basic level, .NET is a platform designed to enable development in multiple languages as long as the application is deployed on Windows. On the other hand, J2EE is designed to enable applications to be deployed on any platform as long as they are written in Java.² Obviously, an organization may choose to develop different application systems using different methodologies and platforms, depending upon deployment and implementation requirements.

Another difference is that .NET is software that can be purchased from Microsoft. J2EE is a set of specifications (managed by Oracle, formerly Sun Microsystems), each of which dictates how various J2EE functions must operate. IBM’s WebSphere Application Server is an example of software that implements J2EE. Oracle makes money from J2EE not only by selling J2EE software, but also by licensing the J2EE specifications to independent software vendors, which then implement software according to the specs.

The bottom line is that both .NET and J2EE can be used to build Web services. A Web service is an application that accepts requests from other systems across the Internet (or an intranet), mediated by lightweight, vendor-neutral communications technologies. Web services enable applications to share data and invoke capabilities from other applications without knowledge of how those other applications were built, what operating system or platform they run on, and what devices are used to access them.

Ruby on Rails

Ruby on Rails is an open-source Web application framework for the Ruby programming language. Ruby on Rails includes tools that make common development tasks easier and is an additional application development framework, similar to but separate from both .NET and J2EE.

Truly, an entire book could be dedicated to the different framework

options available for modern application development. But this is a book about database administration, so the cursory overview provided in this section should be sufficient.

Object Orientation and SQL

Many organizations have adopted *object-oriented* (OO) programming standards and languages because of the claimed advantages of the OO development paradigm. The primary advantages of object orientation are faster program development time and reduced maintenance costs, resulting in a better ROI. Piecing together reusable objects and defining new objects based on similar object classes can dramatically reduce development time and costs.

Object orientation can result in a better ROI.

With benefits like these, it is no wonder that object-oriented programming and development is being embraced by many IT organizations. Historically, one of the biggest problems faced by IT is a large project backlog. In many cases, end users are forced to wait for long periods of time for new applications because the backlog is so great and the talent needed to tackle so many new projects is not available. This backlog can sometimes result in some unsavory phenomena such as business people attempting to build their own applications or purchasing third-party packaged applications (and all of the potential administrative burdens that packages carry). So, it is very clear why the OO siren song lures organizations.

However, because OO and relational databases are not inherently compatible, OO programmers tend to resist using SQL. The set-based nature of SQL is not simple to master and is anathema to the OO techniques practiced by Java and C++ developers. All too often insufficient consideration has been given to the manner in which data is accessed, resulting in poor design and faulty performance.

Object orientation is indeed a political nightmare for those schooled in relational tenets. All too often organizations experience political struggles between OO programming proponents and the data resource management group. The OO crowd espouses programs and components as the center of the universe; the data crowd adheres to normalized, shared data with the RDBMS as the center of the universe.

Thanks to the hype surrounding object orientation, the OO crowd may win many of these battles, but data-centered thinking will eventually win the war. The use of data normalization and shared databases to reduce redundancy provides far too many benefits in the long run for it ever to be abandoned. Data has an existence independent of process, and the OO way of encapsulating data within methods masks data from other processes and is inefficient for sharing data across an organization. If the focus shifts away from data management and sound relational practices, data quality will deteriorate and productivity will decline.

If an OO programming language is to be used against a relational database, you will need to marry the relational world to the OO world. This means your applications will not be object oriented in the true sense of the word because the data will not be encapsulated within the method (that is, the program).

You will need to marry the relational world to the OO world.

There are several techniques that can be used to enable an OO programming language to work with a relational database. Serialization, saving data using a flat file representation of the object, is one approach. Because it can be slow and difficult to use across applications, serialization is not commonly used for persisting object data. A second approach is to use XML, which can be stored natively in many relational database systems. However, XML adds a layer of complexity and requires an additional programming skill set. XML is discussed in more detail later in this chapter.

Probably the most common technique is to deploy an ORM (object-relational mapping) solution. Through ORM, an object's attributes are stored in one or more columns of a relational table. Hibernate is a popular ORM library for Java. NHibernate is an adaptation of Hibernate for the .NET framework. Both Hibernate and NHibernate provide capabilities for mapping objects to a relational database by replacing direct persistence-related database accesses with high-level object-handling functions. Another option is Microsoft LINQ, which stands for Language Integrated Query. LINQ provides a set of .NET framework and language extensions for object-relational mapping.

One additional word of caution here: Many people believe that object-relational databases resolve all of these issues. But an object-relational database is not truly object oriented. The term *object-relational* means

basically that the DBMS supports large multimedia data types and gives users the ability to define their own data types and functions—all good things, but not object orientation. So don't get confused over the similarity of the terms.

Types of SQL

SQL, although a single language, comprises multiple types that exhibit different behavioral characteristics and require different development and administration techniques. SQL can be broken down into categories based on execution type, program requirement, and dynamism.

- *SQL can be planned or unplanned.* A planned SQL request is typically embedded into an application program, but it might also exist in a query or reporting tool. At any rate, a planned SQL request is designed and tested for accuracy and efficiency before it is run in a production system. Contrast this with the characteristics of an unplanned SQL request. Unplanned SQL, also called *ad hoc SQL*, is created “on the fly” by end users during the course of business. Most ad hoc queries are created to examine data for patterns and trends that impact business. Unplanned, ad hoc SQL requests can be a significant source of inefficiency and are difficult to tune. How do you tune requests that are constantly written, rewritten, and changed?

SQL can be planned or unplanned, embedded in a program or stand-alone, dynamic or static.

- *SQL can either be embedded in a program or issued stand-alone.* Embedded SQL is contained within an application program, whereas stand-alone SQL is run by itself or within a query, reporting, or OLAP tool.
- *SQL can be dynamic or static.* A dynamic SQL statement is optimized at run time. Depending on the DBMS, a dynamic SQL statement may also be changed at run time. Static SQL, on the other hand, is optimized prior to execution and cannot change without reprogramming. Favor static SQL to minimize the possibility of SQL injection attacks (see the sidebar “[SQL Injection](#)”).

Programmers must be able to quantify each SQL statement being developed in terms of these three qualities. Every SQL statement exhibits one of these properties for each criterion. For example, a certain SQL statement

can be a planned, embedded, static request, or it could be an unplanned, stand-alone, dynamic request. Be sure to use the right type of SQL for the right situation. [Table 5.2](#) outlines situations and the type of SQL that is most useful for that situation. Of course, the information in this table is meant to serve as a broad suggestion only. You should use your knowledge of your environment and the requirements of the user request to arrive at the correct type of SQL solution.

Table 5.2. SQL Usage Considerations

Situation	Execution Type	Program Requirement	Dynamism
Columns and predicates of the SQL statement can change during execution	Planned	Embedded	Dynamic
SQL formulation does not change	Planned	Embedded	Static
Highly concurrent, high-performance transactions	Planned	Embedded	Dynamic or static
Ad hoc one-off queries	Unplanned	Stand-alone	Dynamic
Repeated analytical queries	Planned	Embedded or stand-alone	Dynamic or static
Quick one-time “fix” programs	Unplanned	Embedded or stand-alone	Dynamic or static

SQL Injection

A common form of hack against Web-exposed applications using SQL to access data is the SQL injection attack. Properly designed applications can thwart SQL injection attacks.

A SQL injection attack inserts SQL statements in the fields of a Web form to cause a poorly designed Web application to expose database content to the attacker. Refer to [Chapter 14](#), “[Database Security](#),” for an in-depth discussion of SQL injection and how to combat it.

SQL Coding for Performance

Developing database application programs requires a good amount of effort to ensure that SQL requests are properly coded for performance. A solid understanding of SQL syntax, database structures, and the programming language is imperative. Let’s concentrate first on SQL.

One of the first rules to learn as a database developer is to let SQL, rather

than the program logic, do the work. It is much better to filter out unwanted data at the DBMS level than to do so within the program. You'll achieve better efficiency by avoiding the actual movement of data between the DBMS and the program. For example, it is better to add more WHERE clauses to SQL SELECT statements than to simply select all rows and filter the data programmatically.

Let SQL do the work.

To use another example, consider the cost of a multitable join statement. It will be easier to tune, say, a four-table join for efficiency than four independent SQL SELECT statements that are filtered and joined using application logic. Of course, this assumes an optimal physical database and the possibility of having to tweak that design (such as by adding indexes).

The more work the DBMS can do to filter data, the greater the efficiency should be, because less data will need to be moved between the DBMS and the application program as it runs. Of course, there is much more to optimizing and tuning SQL than this short discussion of the matter. More details are covered in [Chapter 12](#), "[Application Performance](#)."

Querying XML Data

These days, not all data stored in the database will be relational. XML is gaining popularity for persisting complex data and is frequently used in Web-enabled applications and as a means of data transmission. All of the leading DBMS products provide means of storing and managing XML data.³

XML stands for Extensible Markup Language. You may be familiar with HTML, the markup language used to create Web pages. Like HTML, XML is based upon SGML (Standard Generalized Markup Language). SGML is a language for defining markup languages; it was developed and standardized by the International Organization for Standardization (ISO).

XML uses tags to describe the data itself.

Whereas HTML uses tags to describe *how* data appears on a Web page, XML is designed to transport and store data. In other words, XML is somewhat self-describing. It uses tags to describe the *what*, that is, the data itself. The simple syntax of XML makes it easy to process by machine while remaining understandable to people. Once again, let's use HTML as a

metaphor to help us understand XML. HTML uses tags to describe the appearance of data on a page. For example the tag “text” would specify that the “text” data should appear in boldface. XML uses tags to describe the data itself, instead of its appearance. For example, consider the following XML describing a customer address:

[Click here to view code image](#)

```
<CUSTOMER>
  <first_name>Craig</first_name>
  <middle_initial>S.</middle_initial>
  <last_name>Mullins</last_name>
  <company_name>Mullins Consulting, Inc.
</company_name>
  <street_address>15 Coventry Ct.</street_address>
  <city>Sugar Land</city>
  <state>TX</state>
  <zip_code>77479</zip_code>
  <country>USA</country>
</CUSTOMER>
```

XML is actually a meta-language—that is, a language for defining other markup languages. These languages are collected in dictionaries called Document Type Definitions (DTDs). The DTD stores definitions of tags for specific industries or fields of knowledge. Instead of a DTD, an XML schema could be employed for the same purpose.

When data is stored as XML, it cannot be accessed using standard SQL. Instead, it requires either XQuery or SQL/XML extensions.

XQuery

XQuery is a query and programming language designed to query collections of XML data. XQuery uses XPath expression syntax to address specific parts of an XML document. It supplements this with a SQL-like “FLWOR expression” for performing joins. A FLWOR expression is constructed from the five clauses after which it is named:

- FOR
- LET
- WHERE
- ORDER BY
- RETURN

The XQuery language is not just for querying; it also allows for new XML documents to be constructed. However, there are no features in XQuery for updating XML documents or databases. It also does not provide full text search capability. Over time, expect these shortcomings to be remedied, though.

In short, XQuery is a programming language that can be used to express arbitrary XML-to-XML data transformations with strong typing and logical/physical data independence.

SQL/XML

SQL/XML is an extension to the SQL standard that specifies SQL-based extensions for using XML in conjunction with SQL. At a high level, it offers an XML data type along with new routines, functions, and XML-to-SQL data type mappings to support accessing and manipulating XML in SQL databases. SQL/XML is developed to be complementary to XQuery.

At the heart of the SQL/XML specification are the functions that allow users to access, modify, and construct XML elements or attributes. Examples of these functions include XMLDOCUMENT (which returns an XML value with a single document node and zero or more nodes as its children), XMLCONCAT (which returns a forest of XML elements generated from a concatenation of two or more elements), XMLELEMENT (which returns an XML element given an element name, an optional collection of attributes, and zero or more arguments that make up the contents of the element), XMLAGG (which returns an XML sequence that contains an item for each non-null value in a set of XML values), among many others.

SQL/XML also defines functions that allow the user to embed XQuery expressions in SQL statements and to convert complex data types. These functions include XMLQUERY and XMLTABLE.

Defining Transactions

A [*transaction*](#) is an atomic unit of work with respect to recovery and consistency. A logical transaction performs a complete business process typically on behalf of an online user. It may consist of several steps and may comprise more than one physical transaction. The results of running a transaction will record the effects of a business process—a complete business process. The data in the database must be correct and proper after the

transaction executes.

A transaction is an atomic unit of work with respect to recovery and consistency.

When all the steps that make up a specific transaction have been accomplished, a COMMIT is issued. The COMMIT signals that all work since the last COMMIT is correct and should be externalized to the database. At any point within the transaction, the decision can be made to stop and roll back the effects of all changes since the last COMMIT. When a transaction is rolled back, the data in the database will be restored to the original state before the transaction was started. The DBMS maintains a transaction log (or journal) to track database changes.

Transactions exhibit ACID properties. ACID is an acronym for [*atomicity*](#), [*consistency*](#), [*isolation*](#), and [*durability*](#). Each of these four qualities is necessary for a transaction to be designed correctly.

- [*Atomicity*](#) means that a transaction must exhibit “all or nothing” behavior. Either all of the instructions within the transaction happen, or none of them happen. Atomicity preserves the “completeness” of the business process.
- [*Consistency*](#) refers to the state of the data both before and after the transaction is executed. A transaction maintains the consistency of the state of the data. In other words, after a transaction is run, all data in the database is “correct.”
- [*Isolation*](#) means that transactions can run at the same time. Any transactions running in parallel have the illusion that there is no concurrency. In other words, it appears that the system is running only a single transaction at a time. No other concurrent transaction has visibility to the uncommitted database modifications made by any other transactions. To achieve isolation, a locking mechanism is required.
- [*Durability*](#) refers to the impact of an outage or failure on a running transaction. A durable transaction will not impact the state of data if the transaction ends abnormally. The data will survive any failures.

Let’s use an example to better understand the importance of transactions to database applications. Consider a banking application. Assume that you wish to withdraw \$50 from your account with Mega Bank. This “business process”

requires a transaction to be executed. You request the money either in person by handing a slip to a bank teller or by using an ATM (automated teller machine). When the bank receives the request, it performs the following tasks, which make up the complete business process. The bank will

1. Check your account to make sure you have the necessary funds to withdraw the requested amount
2. If you do not, deny the request and stop; otherwise continue processing
3. Debit the requested amount from your checking account
4. Produce a receipt for the transaction
5. Deliver the requested amount and the receipt to you

Design transactions that ensure ACID properties.

The transaction that is run to perform the withdrawal must complete all of these steps, or none of these steps, or else one of the parties in the transaction will be dissatisfied. If the bank debits your account but does not give you your money, you will not be satisfied. If the bank gives you the money but does not debit the account, the bank will be unhappy. Only the completion of every one of these steps results in a “complete business process.” Database developers must understand the requisite business processes and design transactions that ensure ACID properties.

To summarize, a transaction—when executed alone, on a consistent database—will either complete, producing correct results, or terminate, with no effect. In either case the resulting condition of the database will be a consistent state.

Transaction Guidelines

A transaction should be short in duration because it locks shared resources. Of course, “short” will vary from system to system. A short transaction in a very large system handling multiple thousands of transactions per second will most likely be measured in subseconds.

A transaction should be short in duration.

At any rate, transactions must be designed to remove the human element of “think time” from the equation. When a transaction locks resources, it makes those resources inaccessible to other transactions. Therefore, a good

transaction must be designed so that it does not wait for user input in the middle of the processing.

Unit of Work

Unit of work (UOW) is another transaction term that describes a physical transaction. A UOW is a series of instructions and messages that, when executed, guarantee data integrity. So a UOW and a transaction are similar in concept. However, a UOW is not necessarily a complete business process—it can be a subset of the business process, and a group of units of work can constitute a single transaction.

A UOW is a series of instructions and messages that guarantee data integrity.

Each UOW must possess ACID characteristics. In other words, if the transaction were to fail, the state of the data upon failure must be consistent in terms of the business requirements.

Transaction Processing Systems

A transaction processing (TP) system, appropriately enough, facilitates the processing of transactions. Such a system is sometimes referred to as a *transaction server* or a *transaction processing monitor*. Regardless of name, a TP system delivers a scheme to monitor and control the execution of transaction programs. The TP system also provides an API—a mechanism for programs to interact and communicate with the TP server. Examples of TP systems include CICS, IMS/TM, Tuxedo, and Microsoft Transaction Server.

A TP system delivers a scheme to monitor and control the execution of transaction programs.

The TP system provides an environment for developing and executing presentation logic and business logic components. A TP system is useful for mission-critical applications requiring a high volume of concurrent users with minimal downtime. Used properly, a TP system can efficiently control the concurrent execution of many application programs serving large numbers of online users. Another major benefit of some TP systems is their ability to ensure ACID properties across multiple, heterogeneous databases. This is accomplished using a two-phase COMMIT, where the TP system controls the issuance of database commits and ensures their satisfactory completion. If

your application requires online access and modification of heterogeneous databases, a TP system is recommended.

A *transaction server* is ideal for building high-performance and reliable distributed applications across heterogeneous environments. A TP system can support the diverse application requirements of front-end e-commerce applications as well as the robust needs of back-office processes. When platform independence is crucial, a TP system can help developers to successfully develop, manage, and deploy online applications that are completely independent of the underlying communications, hardware, and database environment.

Of course, a TP system is not required in order to develop database transactions for every application. The DBMS itself can deliver ACID properties for the data it manages. Yet even in a single-DBMS environment, a TP system can provide development benefits. Take a look at [Figure 5.4](#). It shows the typical application setup: a database server without a TP system. Requests are made by the client through the presentation layer to the database server (DBMS). The client can make data requests directly of the database server, or the client may execute stored procedures to run application logic on the database server.

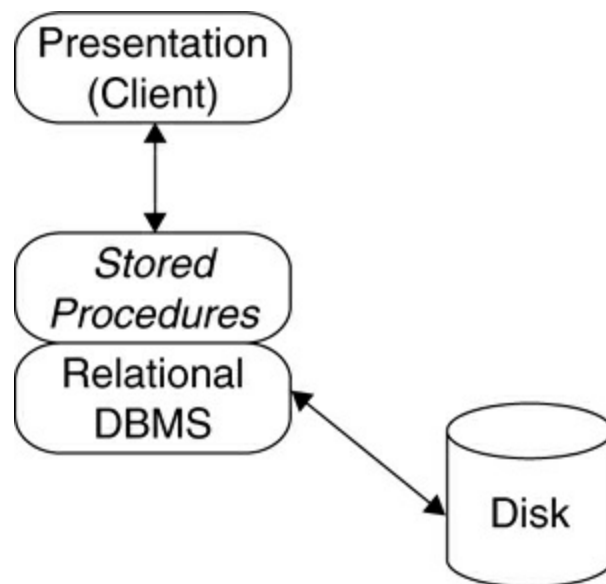


Figure 5.4. Using a database server

Now let's compare [Figure 5.4](#) with [Figure 5.5](#), which shows an application using a TP system. Database requests are run within the TP system, which

helps to control the transaction workflow. Furthermore, the TP system enables the application to make requests of multiple heterogeneous databases and to coordinate database modifications.

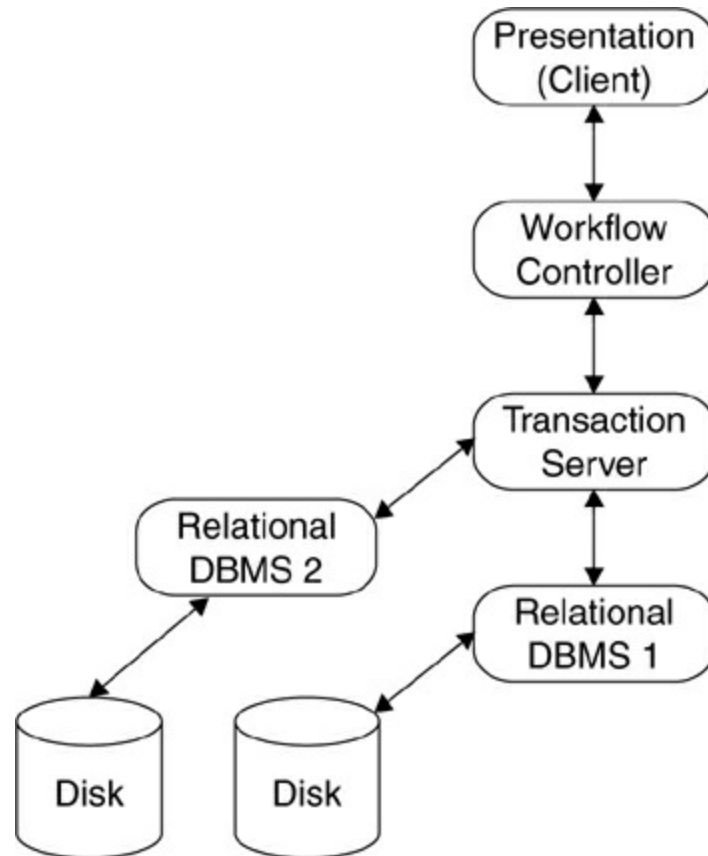


Figure 5.5. Using a transaction server

If your database applications use a TP system, the DBA should work with the system administrator to ensure that the DBMS and the TP system are set up to work with each other in an efficient manner. Furthermore, both the DBA and the SA will need to monitor the interaction between the DBMS and the TP system on an ongoing basis.

Ensure that the DBMS and the TP system are set up to work together in an efficient manner.

Application Servers

A more recent type of middleware for serving database transactions is the *application server*. An application server usually combines the features of a transaction server with additional functionality to assist in building,

managing, and distributing database applications. An application server, such as IBM's WebSphere, provides an environment for developing and integrating components into a secure, performance-oriented application.

Other examples of application servers include Zend Server (for PHP-based applications), Base4 Application Server (open source), and TNAPS Application Server (freeware).

There are several advantages of using an application server. When business logic is centralized on an application server, application updates will apply to all users. The risk of having old versions of the application is eliminated. Furthermore, changes to the application configuration, such as moving a database server, can occur centrally and thereby apply to all users at the same time. Additionally, an application server can act as a central point for application security.

With an application server, application changes are made centrally, thereby reducing risk.

Application servers also serve many of the purposes of a transaction server, such as the delivery of transaction support, whereby a unit of work can be made atomic and indivisible.

Locking

Every programmer who has developed database programs understands the potential for concurrency problems. When one program tries to read data that is in the process of being changed by another program, the DBMS must prohibit access until the modification is complete, in order to ensure data integrity. The DBMS uses a [locking](#) mechanism to enable multiple, concurrent users to access and modify data in the database. By using locks, the DBMS automatically guarantees the integrity of data. The DBMS locking strategies permit multiple users from multiple environments to access and modify data in the database at the same time.

Locks are used to ensure the integrity of data. When a database resource is locked by one process, another process is not permitted to change the locked data. Locking is necessary to enable the DBMS to facilitate the ACID properties of transaction processing.

Locks are used to ensure the integrity of data.

Data may be locked at different levels within the database. It is the DBA's job to determine the appropriate level of locking for each database object, based on how the data will be accessed and to what extent concurrent users will access the data. Theoretically, database locks can be taken at the following levels:

- Column
- Row
- Page (or block)
- Table
- Tablespace
- Database

The level of locking is known as *lock granularity*. The actual lock granularity levels available will depend on the DBMS in use. Typically the lock granularity is controlled by the DBA when the database object is created, but there may be an overall default lock granularity that is used if none is specified. Nevertheless, it is a good practice for DBAs always to specify the lock granularity when database objects are created.

The lock granularity specification should be based on the needs of the applications and users that will be accessing and changing the data. In general, the smaller the granularity of the lock, the more concurrent access will be allowed, as shown in [Figure 5.6](#). However, the smaller the granularity of the lock, the more resources the DBMS will need to consume to perform locking.

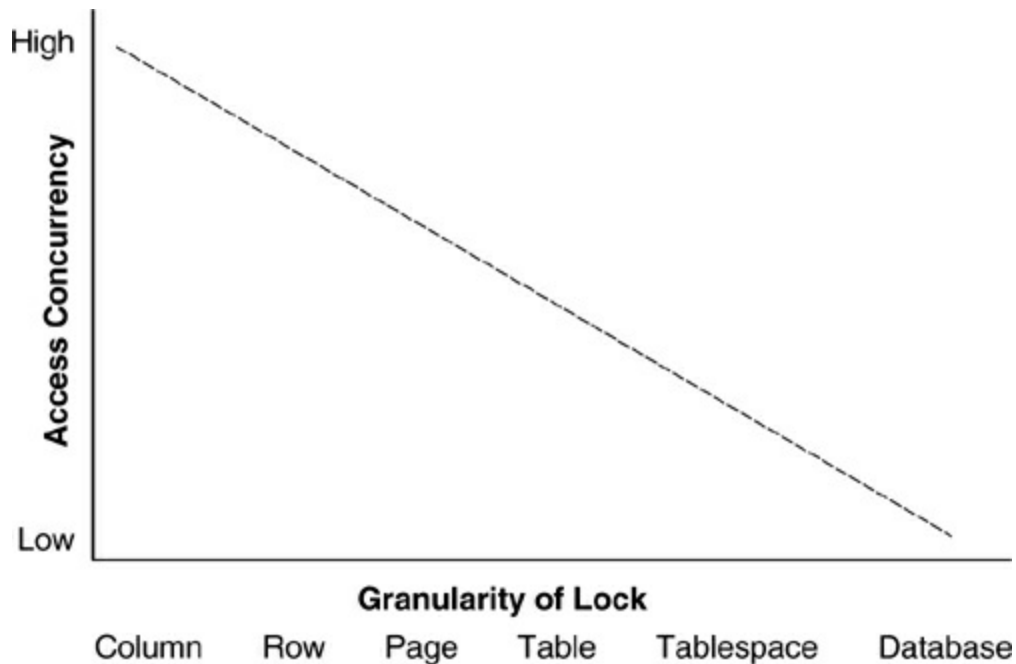


Figure 5.6. Lock granularity and concurrent access

The smallest unit that can conceivably be locked is a single column; the largest lock size is the entire database. Neither of these two options is really practical. Locking individual columns would maximize the ability to concurrently update data, but it would require too much overhead. Conversely, locking the database would be cheap to implement, but it would restrict concurrency to the point of making the database worthless for anybody but a single user.

Most database implementations choose between row and page locking. For many applications, page locking is sufficient. However, applications needing many concurrent update processes may benefit from the smaller granularity of row locking. Very infrequently, when a specific process will be run when no other concurrent access is needed, table locking can be useful. For this reason many DBMS products provide a LOCK TABLE command that can be issued to override the current locking granularity for a database object during the course of a single process or UOW. The DBA must analyze the application processing requirements for each database object to determine the optimal locking granularity.

Most database implementations choose between row and page locking.

Index entries can also be locked, depending on the DBMS and version in use. However, index locking can be a significant impediment to performance. Because index entries are usually quite small, it is not uncommon for locking to block application access—especially when the DBMS locks indexes at the block or page level. Some DBMSs do not require index locks, instead handling integrity by using locks on the data. Remember, there is no data in the index that is not also in a table.

Index locking can be a significant impediment to performance.

The exact nature of locking and the types of locks taken will differ from DBMS to DBMS. This section will cover the basics of locking that are generally applicable to most DBMS products.

Types of Locks

At a very basic level, a DBMS will take a *write lock* when it writes information or a *read lock* when it reads information. A write occurs for INSERT, UPDATE, and DELETE statements. A read occurs for SELECT statements. But to actually accomplish such locking, the typical DBMS will use three basic types of locks: shared locks, exclusive locks, and update locks.

- A [*shared lock*](#) is taken by the DBMS when data is read with no intent to update it. If a shared lock has been taken on a row, page, or table, other processes or users are permitted to read the same data. In other words, multiple processes or users can have a shared lock on the same data.
- An [*exclusive lock*](#) is taken by the DBMS when data is modified. If an exclusive lock has been taken on a row, page, or table, other processes or users are generally not permitted to read or modify the same data. In other words, multiple processes or users cannot have an exclusive lock on the same data.
- An [*update lock*](#) is taken by the DBMS when data must first be read before it is changed or deleted. The update lock indicates that the data may be modified or deleted in the future. If the data is actually modified or deleted, the DBMS will promote the update lock to an exclusive lock. If an update lock has been taken on a row, page, or table, other processes or users generally are permitted to read the data, but not to modify it. So a single process or user can have an update lock while

other processes and users have shared locks on the same data. However, multiple processes or users cannot have both an exclusive lock and an update lock, or multiple update locks, on the same data.

Intent Locks

In addition to shared, exclusive, and update locks, the DBMS also will take another type of lock, known as an *intent lock*. Intent locks are placed on higher-level database objects when a user or process takes locks on the data pages or rows. An intent lock stays in place for the life of the lower-level locks.

An intent lock stays in place for the life of the lower-level locks.

For example, consider a table created with row-level locking. When a process locks the row, an intent lock is taken on the table. Intent locks are used primarily to ensure that one process cannot take locks on a table, or pages in the table, that would conflict with the locking of another process. For example, if a user was holding an exclusive row lock and another user wished to take out an exclusive table lock on the table containing the row, the intent lock held on the table by the first user would ensure that its row lock would not be overlooked by the lock manager.

Lock Time-outs

When data is locked by one process, other processes must wait for the lock to be released before processing the data. A lock that is held for a long time has the potential to severely degrade performance because other processes must wait until the lock is released and the data becomes available. Furthermore, if the application is designed improperly or has a bug, the blocking lock may not be released until the program fails or the DBA intervenes.

The locking mechanism of the DBMS prevents processes from waiting forever for a lock to be released by timing out. Each DBMS provides a parameter to set a *lock time-out value*. Depending on the DBMS, the lock time-out value might be set at the DBMS level, the process level, or the connection level. Regardless of the level, after a process waits for the predetermined amount of time for a lock to be granted, the process will receive an error message informing it that the time-out period has been exceeded. Such an approach assumes that a problem occurs after a certain

amount of time is spent waiting for a lock. Time-outs prevent a process from waiting indefinitely for locks—the rationale being that it is better for a process to give up and release its locks than to continue to wait and perhaps block other processes from running.

It is usually a good practice for programs to retry an operation when a lock time-out error is received. If multiple lock time-outs occur for the same operation, the program should log the problem and inform the user that it cannot proceed.

Deadlocks

Another locking problem that can occur is deadlocking. A [deadlock](#) occurs when concurrent processes are competing for locks. [Figure 5.7](#) shows a deadlock situation. Process A holds a lock on row 3 and is requesting a lock on row 7; process B holds a lock on row 7 and is requesting a lock on row 3.




Figure 5.7. A deadlock situation

A deadlock occurs when concurrent processes are competing for locks.

The deadlock is a specific type of lock time-out. It occurs when one process holds a lock that another process is requesting at the same time the second process holds a lock that the first process is requesting. This is also known as a “deadly embrace.” The DBMS will choose one of the processes to abort and roll back so the other process can continue.

As with time-outs, it is a good practice for programs to retry an operation when a deadlock error is received. If multiple deadlocks occur for the same operation, the program should log the problem and inform the user that it cannot proceed.

If deadlocks are a persistent problem, application or database design changes may be warranted. One technique is to consider changing the lock granularity (perhaps from page to row) so that less data is locked for each lock request. Another technique is to make application changes. You might rewrite the program to lock all required resources at the beginning of the transaction. However, such a technique is not likely to be possible for most applications. A final method of avoiding deadlocks is to ensure that all database modifications occur in the same order for all programs in your shop. If the updates in every program are sequenced, deadlocks can be minimized or eliminated. Perhaps you'll find a logical business order for updates that is reasonable to use for applications. If you don't, consider using an arbitrary policy such as modification in alphabetical order.

Lock Duration

Lock duration refers to the length of time that a lock is held by the DBMS. The longer a lock is held on a database resource, the longer a concurrent lock request must wait to be taken. As lock durations increase, so does the likelihood of lockout time-outs.

Lock duration refers to the length of time that a lock is held by the DBMS.

Each DBMS provides parameters that can be set to impact the duration of a lock. Typically these parameters are set at the program, transaction, or SQL statement level. In general, there are two parameters that affect lock duration: isolation level and acquire/release specification.

Isolation Level

The *isolation level* specifies the locking behavior for a transaction or statement. Standard SQL defines four isolation levels that can be set using the SET TRANSACTION ISOLATION LEVEL statement:

- UNCOMMITTED READ
- COMMITTED READ

- REPEATABLE READ
- SERIALIZABLE

The preceding list progresses from lowest to highest isolation. The higher the isolation level, the stricter the locking protocol becomes. The lower the isolation level, the shorter the lock duration will be. Additionally, each higher isolation level is a superset of the lower levels. Let's briefly examine each of the standard isolation levels.

The isolation level specifies the locking behavior for a transaction or statement.

Specifying *UNCOMMITTED READ* isolation implements read-through locks and is sometimes referred to as a *dirty read*. It applies to read operations only. With this isolation level, data may be read that never actually exists in the database, because the transaction can read data that has been changed by another process but is not yet committed. *UNCOMMITTED READ* isolation provides the highest-level availability and concurrency of the isolation levels, but the worst degree of data integrity. It should be used only when data integrity problems can be tolerated. Certain types of applications, such as those using analytical queries, estimates, and averages, are often candidates for *UNCOMMITTED READ* locking. A dirty read can cause duplicate rows to be returned where none exist, or no rows to be returned when one (or more) actually exists. When choosing *UNCOMMITTED READ* isolation, the programmer and DBA must ensure that these types of problems are acceptable for the application. Refer to the sidebar "[Dirty Read Scenarios](#)" for additional guidance on when to consider using *UNCOMMITTED READ* isolation.

COMMITTED READ isolation, also called *cursor stability*, provides more integrity than *UNCOMMITTED READ* isolation. When *READ COMMITTED* isolation is specified, the transaction will never read data that is not yet committed; it will only *COMMITTED READ* data.

REPEATABLE READ isolation places a further restriction on reads, namely, the assurance that the same data can be accessed multiple times during the course of the transaction without its value changing. The lower isolation levels (*UNCOMMITTED READ* and *COMMITTED READ*) permit the underlying data to change if it is accessed more than once. Use *REPEATABLE READ* isolation only when data can be read multiple times

during the course of the transaction and the data values must be consistent.

Dirty Read Scenarios

When is using UNCOMMITTED READ isolation appropriate? The general recommendation is to avoid it if your results must be 100 percent accurate. For example, avoid dirty reads when calculations must balance, when data is being retrieved to insert into another table, or for mission-critical data where consistency and integrity are crucial.

Frankly, most database applications are not candidates for dirty reads. However, there are specific situations where permitting uncommitted data to be read is beneficial.

One such situation is when accessing a reference, code, or lookup table that is very static. Because the data is not volatile, an UNCOMMITTED READ would usually be no different from using a stricter isolation level. For those rare occasions when the lookup codes are being modified, the problems should be minimal for concurrent transactions.

When a transaction must perform statistical processing on a large amount of data, a dirty read may be useful. For example, consider a transaction designed to return the average account balance for each type of account. The impact of using UNCOMMITTED READ isolation is minimal because changing a single value should not have a significant impact on the result. Because the result is an average of multiple values, one or perhaps a few “bad” values are unlikely to change the average significantly.

Data warehousing queries are good candidates for dirty reads. A data warehouse is a time-sensitive, subject-oriented store of business data that is used for analytical processing. Other than periodic data propagation and/or replication, access to the data warehouse is read only. An UNCOMMITTED READ can cause little damage because the data generally does not change.

You might also consider using dirty read for those rare instances when a transaction accesses a table, or set of tables, that is used by a single user only. If only one person can modify the data, the application programs can be coded so that most reads can use UNCOMMITTED

READ isolation with no negative impact on data integrity.

Finally, if the data being accessed is already inconsistent, little harm can be done by using a dirty read to access the information.

Finally, *SERIALIZABLE* isolation provides the greatest integrity. *SERIALIZABLE* isolation removes the possibility of phantoms. A phantom occurs when the transaction opens a cursor that retrieves data and another process subsequently inserts a value that would satisfy the request and should be in the results set. For example, consider the following situation:

- Transaction 1 opens a cursor and reads account information, keeping a running sum of the total balance for the selected accounts.
- Transaction 2 inserts a new account that falls within the range of accounts being processed by Transaction 1, but the insert occurs after Transaction 1 has passed the new account.
- Transaction 2 COMMITs the insert.
- Transaction 1 runs a query to sum the values to check the accuracy of the running total. However, the totals will not match.

SERIALIZABLE isolation provides the greatest integrity.

SERIALIZABLE isolation eliminates this problem. Phantoms can occur for lower isolation levels, but not when the isolation level is *SERIALIZABLE*.

Most DBMS products support the specification of isolation level at the program or transaction level, as well as at the SQL statement level.

Keep in mind that your DBMS may not implement all of these isolation levels, or it may refer to them by other names. Be sure you understand the isolation supported by each DBMS you use and its impact on application behavior and lock duration.

Acquire/Release Specification

An additional parameter that impacts lock duration is the treatment of intent locks. Regular transaction locks are taken as data is accessed and modified. However, some DBMS products provide methods to control when intent locks are taken. Intent locks can be acquired either immediately when the transaction is requested or iteratively as needed while the transaction

executes. Furthermore, intent locks can be released when the transaction completes or when each intent lock is no longer required for a unit of work.

If the DBMS supports different options for the acquisition and release of intent locks, the parameter will be specified at the transaction or program level.

Skipping Locked Rows

An additional locking option, available in some DBMSs,⁴ is the ability to skip locked data. If you code a parameter specifying SKIP LOCKED DATA on certain SQL statements, any data that is locked will simply be skipped over instead of the DBMS waiting for the lock to be released.

This option should be used only sparingly and only with a full understanding of its impact. When you tell the DBMS to skip locked data, that data is not accessed and will not be available to your program. The benefit, of course, is improved performance because you will not incur any lock wait time; however, it comes at the cost of not accessing the locked data.

You should use this feature sparingly and with extreme caution. Before skipping locked data, make sure that you completely understand exactly what you are telling the DBMS to do. It is very easy to misuse this feature and wind up reading less data than you want.

Consider using this option in certain test environments, in a data warehouse where data is read only, and possibly even in production under the proper conditions. For example, perhaps you have a program that needs to read from a table such as a queue to get a next number. If it is not imperative that the numbers be sequential, skipping locked data can eliminate bottlenecks by skipping any locked rows/pages to get data off of the queue.

Lock Escalation

If processing causes the DBMS to hold too many locks, lock escalation can occur. [*Lock escalation*](#) is the process of increasing the lock granularity for a process or program. When locks are escalated, more data is locked, but fewer actual locks are required. An example of escalating a lock would be moving from page locks to table locks. You can see where this would minimize the number of locks the DBMS needs to track—multiple page locks for a table can be converted into a single lock on the entire table. Of course, this impacts concurrent access because the process locks the entire table, making it

inaccessible to other processes.

Lock escalation is the process of increasing the lock granularity for a process or program.

The DBMS kicks off lock escalation based on preset thresholds. Typically, the DBMS will provide system parameters that can be set to customize the actual manner in which the DBMS escalates locks, or to turn off lock escalation. Also, the DBMS will provide DDL parameters for database objects to indicate on an object-by-object basis whether escalation should occur.

Some DBMSs, such as DB2 and Microsoft SQL Server, provide the capability to escalate locks, whereas others, such as Oracle, do not. However, both DB2 and Microsoft SQL Server can escalate from page locks to table locks or from row locks to table locks. Neither allows escalation from row locks to table locks.

Programming Techniques to Minimize Locking Problems

We have learned that locking is required to ensure data integrity. If application programs are not designed with database locking in mind, though, problems can arise. Application developers must understand the impact of locking on the performance and availability of their applications. If locks are held too long, time-outs will make data less available. If applications request locks in a disorganized manner, deadlocks can occur, causing further availability problems.

Standardize the sequence of updates within all programs.

Fortunately, though, there are development techniques that can be applied to minimize locking problems. One such technique is to standardize the sequence of updates within all programs. When the sequence of updates is the same for all programs, deadlocks should not occur.

Another programming technique is to save all data modification requests until the end of the transaction. The later modifications occur in a transaction, the shorter the lock duration will be. From a logical perspective, it really does not matter where a modification occurs within a transaction, as long as all of the appropriate modifications occur within the same transaction. However, most developers feel more comfortable scattering the data modification logic

throughout the transaction in a pattern that matches their concept of the processes in the transaction. Grouping modifications such as INSERT, UPDATE, and DELETE statements and issuing them near the end of the transaction can improve concurrency because resources are locked for shorter durations.

Locking Summary

Database locking is a complex subject with more details than we have covered in this section. Each DBMS performs locking differently, and you will need to study the behavior of each DBMS you use to determine how best to set locking granularity and isolation levels, and to program to minimize time-outs and deadlocks.

Batch Processing

Most of the discussion in this chapter has centered around transaction processing, which is usually assumed to be for online processes. However, many database programs are designed to run as batch jobs with no online interaction required. DBAs must be aware of the special needs of *batch database programs*.

The first design concern for batch database programs is to ensure that database COMMITs are issued within the program. Except for very trivial programs that access small amounts of data, database COMMITs should be issued periodically within a batch program to release locks. Failure to do so can cause problems such as a reduction in availability for concurrent programs because a large number of locks are being held, or a large disruption if the batch program aborts, because all the database modifications must be rolled back.

Ensure that database COMMITs are issued within the program.

Additionally, if a batch program with no COMMITs fails, all of the work that is rolled back must be performed again when the problem is resolved and the batch program is resubmitted for processing. A batch program with COMMITs must be designed for restartability: The batch program must keep track of its progress by recording the last successful COMMIT and including logic to reposition all cursors to that point in the program. When the program

is restarted, it must check to see if it needs to reposition and, if so, execute the repositioning logic before progressing.

Another problem that occurs frequently with batch database program development is a tendency for developers to think in terms of file processing, rather than database processing. This is especially true for mainframe COBOL programmers who have never worked with database systems. Each developer must be trained in the skills of database programming, including SQL skills, set-at-a-time processing, and database optimization. The responsibility for assuring that developers have these skills quite often falls on the DBA.

Finally, batch programs typically are scheduled to run at predetermined times. The DBA should assist in batch database job scheduling to help minimize the load on the DBMS. Batch jobs that are long running and resource consuming should be scheduled during off-peak online transaction processing hours.

Summary

Application design and development is the job of systems analysts and application programmers. However, the DBA must be involved in the process when programs are being written to access databases. Special skills are required that can be difficult to master. The DBA must first understand these skills and then work to transfer the knowledge to developers. This is a continual job because new programmers are constantly being hired—each with a different level of skill and degree of database experience. Furthermore, DBMS products are constantly changing, resulting in additional development options and features that need to be mastered.

Review

1. Describe what the acronym [*ACID*](#) means and define each component.
2. What is ORM and why would it be necessary for application development?
3. Why is locking required to assure data integrity?
4. Describe the difference between CURSOR STABILITY and REPEATABLE READ isolation levels.
5. Under what circumstance should an isolation level of

UNCOMMITTED READ be considered?

6. Describe two application design techniques to minimize the impact of locking on application performance.
7. What does *relational closure* mean, and what is its significance in application design?
8. Describe, at a high level, what is required to embed SQL into an application program written in a programming language like C or Visual Basic.
9. What is the difference between a lock time-out and a deadlock?
10. What programming techniques can be used to minimize deadlocks and why?

Bonus Question

Why might the order of database modifications within a transaction impact deadlocks?

Suggested Reading

Anagol-Subbarao, Anjali. *J2EE Web Services on BEA WebLogic*. Upper Saddle River, NJ: Prentice Hall (2005). ISBN 0-13-143072-6.

Applequist, Daniel K. *XML and SQL: Developing Web Applications*. Boston, MA: Addison-Wesley (2002). ISBN 0-201-65796-1.

Bales, Donald. *Java Programming with Oracle JDBC*. Sebastopol, CA: O'Reilly (2002). ISBN 0-596-00088-X.

Barcia, Roland, et al. *Persistence in the Enterprise: A Guide to Persistence Technologies*. Upper Saddle River, NJ: IBM Press (2008). ISBN 978-0-13-158756-4.

Beighley, Lynn, and Michael Morrison. *Head First PHP & MySQL*. Sebastopol, CA: O'Reilly (2009). ISBN 978-0-596-00630-3.

Bernstein, Philip A., and Eric Newcomer. *Principles of Transaction Processing*. San Francisco, CA: Morgan Kaufmann (1997). ISBN 1-55860-415-4.

Ceri, Stefano, et al. *Designing Data-Intensive Web Applications*. San Francisco, CA: Morgan Kaufmann (2003). ISBN 1-55860-190-2.

Date, C. J., with Hugh Darwen. *A Guide to the SQL Standard*. 4th ed.

- Reading, MA: Addison-Wesley (1997). ISBN 0-201-96426-0.
- Dix, Paul. *Service-Oriented Design with Ruby and Rails*. Boston, MA: Addison-Wesley (2010). ISBN 978-0-321-65936-1.
- Donahoo, Michael J., and Gregory D. Speegle. *SQL: Practical Guide for Developers*. San Francisco, CA: Morgan Kaufmann (2005). ISBN 978-0-12-220531-6.
- Fronckowiak, John W. *Teach Yourself OLE DB and ADO in 21 Days*. Indianapolis, IN: SAMS Publishing (1997). ISBN 0-672-31083-X.
- Garvin, Curtis, and Steve Eckols. *DB2 for the COBOL Programmer Part 1*. 2nd ed. Fresno, CA: Mike Murach & Associates (1999). ISBN 1-890774-16-2.
- Garvin, Curtis, and Anne Prince. *DB2 for the COBOL Programmer Part 2*. 2nd ed. Fresno, CA: Mike Murach & Associates (1999). ISBN 1-890774-03-0.
- Gray, Jim, and Andreas Reuter. *Transaction Processing: Concepts and Techniques*. San Francisco, CA: Morgan Kaufmann (1993). ISBN 1-55860-190-2.
- Geiger, Kyle. *Inside ODBC*. Redmond, WA: Microsoft Press (1995). ISBN 1-55615-815-7.
- Gulutzan, Peter, and Trudy Pelzer. *SQL-99 Complete, Really*. Lawrence, KS: R&D Books (1999). ISBN 0-87930-568-1.
- Harrington, Jan L. *SQL Clearly Explained*. 3rd ed. Burlington, MA: Morgan Kaufmann (2010). ISBN 978-0-12-375697-8.
- Jennings, Roger. *Database Developer's Guide with Visual Basic 6*. Indianapolis, IN: SAMS Publishing (1999). ISBN 0-672-31063-5.
- Jepson, Brian. *Java Database Programming*. New York, NY: John Wiley & Sons (1997). ISBN 0-471-16518-2.
- Kaute, Pierre Henri, Tobin Harris, Christian Bauer, and Gavin King. *NHibernate in Action*. Greenwich, CT: Manning Publications (2009). ISBN 978-1-932394-92-4.
- Kline, Kevin, with Daniel Kline. *SQL in a Nutshell*. 3rd ed. Sebastopol, CA: O'Reilly (2009). ISBN 978-0-596-51884-4.
- Lewis, Philip M., Arthur Bernstein, and Michael Kifer. *Databases and*

- Transaction Processing*. Boston, MA: Addison-Wesley (2002). ISBN 0-201-70872-8.
- Loosley, Chris, and Frank Douglas. *High-Performance Client/Server*. New York, NY: John Wiley & Sons (1998). ISBN 0-471-16269-8.
- Marguerie, Fabrice, Steve Eichert, and Jim Wooley. *LINQ in Action*. Greenwich, CT: Manning Publications (2008). ISBN 978-1-933988-16-0.
- Melton, Jim. *Understanding SQL's Stored Procedures: A Complete Guide to SQL/PSM*. San Francisco, CA: Morgan Kaufmann (1998). ISBN 1-55860-461-8.
- Patrick, Tim. *ADO.NET 4: Step by Step*. Sebastopol, CA: Microsoft Press (2010). ISBN 978-0-7356-3888-4.
- Price, Jason. *Java Programming with Oracle SQLJ*. Sebastopol, CA: O'Reilly (2001). ISBN 0-596-00087-1.
- Pugh, Eric, and Joseph D. Gradecki. *Professional Hibernate*. Indianapolis, IN: Wrox (2004). ISBN 0-7645-7677-1.
- Reese, George. *Java Database Best Practices*. Sebastopol, CA: O'Reilly (2003). ISBN 0-596-00522-9.
- Rinehart, Martin. *Java Database Development*. Berkeley, CA: McGraw-Hill (1998). ISBN 0-07-882356-0.
- Sceppa, David. *Programming Microsoft ADO.NET*. Sebastopol, CA: Microsoft Press (2012). ISBN 978-0-7356-4801-2.
- Syverson, Bryan. *Murach's SQL for SQL Server*. Fresno, CA: Mike Murach & Associates (2002). ISBN 1-890774-16-2.
- Walmsley, Priscilla. [XQuery](#). Sebastopol, CA: O'Reilly (2007). ISBN 978-0-596-00634-1.
- Yank, Kevin. *Build Your Own Database Driven Web Site Using PHP and MySQL*. 3rd ed. Collingwood, VIC, Australia: SitePoint (2004). ISBN 0-9752402-1-8.

6. Design Reviews

Decisions made during the database design process and the application development life cycle (ADLC) must be reviewed to ensure correctness. This is the purpose of a design review.

What Is a Design Review?

Design reviews are an important facet of the ADLC for database applications. It is during the design review that all aspects of the database and application code are reviewed for efficiency, effectiveness, and accuracy. It is imperative that all database applications, regardless of their size, be reviewed to ensure that the application is designed properly, that the coding techniques are cost-effective, and that the database can be accessed and modified correctly and efficiently. The design review is an important process for checking the validity of design decisions and correcting errors before applications and databases are promoted to production status. Design reviews also can be used to ensure that developers have adhered to corporate standards for internal documentation, naming schemes, maintainability, and so on.

All aspects of the database and application code are reviewed for efficiency, effectiveness, and accuracy.

Multiple design reviews should be conducted over the course of an application's life. For database applications, the DBA must participate in every design review, at every stage. It is imperative that the application be reviewed before, during, and after implementation. Design reviews are critical for ensuring that an application is properly designed to achieve its purpose.

Design reviews address many aspects of the development process and its resulting application. Imposing the design review process on an application exposes it to a thorough review of every underlying component, structure, and nuance of the application. Some of the areas that can be addressed by a design review include

- A validation of the intent and purpose of the application
- An assessment of the logical data model

- An assessment of the physical data model
- A review and analysis of the physical DBMS parameters
- A prediction of SQL performance
- A judgment on the practicality of the programming language techniques deployed
- An analysis of overall performance after production implementation

A group consisting of subject matter experts and the developer's peers and coworkers should conduct each design review. The DBA usually must act as the focal point for organizing and conducting design reviews. Frankly, if the DBA does not organize design reviews, it is unlikely that any design review will be conducted. In addition, if design reviews are not conducted, the application is more apt to suffer performance and availability problems in a production environment.

Subject matter experts, peers, and coworkers should conduct each design review.

Rules of Engagement

Let's cover the ground rules of a design review before defining each type of review. A design review is conducted by a group of people—each having different backgrounds, skills, expectations, and opinions. When any such group is convened to discuss potential problems and errors, confrontation is unavoidable. Each participant must possess the ability to discuss and reach consensus on issues without turning the review into an unproductive battle or argument. To accomplish this, participants must avoid being combative. Everyone needs to understand that the only goal is to promote the best-performing, most usable application possible.

One of the biggest threats to the success of a design review is the possibility that negative criticism will be perceived as a personal attack. If the atmosphere of the review is threatening or the developer perceives it to be so, the developer is likely to resist contributing to the review or accepting an unbiased critique of the work. To avoid such a scenario, be sure that all participants back up their assertions and suggestions with facts, manual references, white papers, articles, blog citations, and experience. As much as possible, participants should check their emotions at the door.

All participants should back up their assertions and suggestions with facts, manual references, white papers, articles, blog citations, and experience.

Design Review Participants

With respect to choosing the design review participants, two guidelines will help to ensure success. The first is to create formal roles for the design review and assign the proper individuals to fulfill those roles. The second is to make sure that participants possess the appropriate skills to actively engage in the design review process.

Create formal roles for the design review.

First, let's discuss the different roles required for a successful design review. Formal roles should include a leader, a scribe, a mediator, and the participants.

The Leader

It is imperative that each design review have only one leader. The leader can change from one design review to the next, but within the scope of a single design review, a single leader must be assigned. The leader's role is multifaceted. The leader

- Acts as a master of ceremonies to keep the review process moving along
- Creates and follows an agenda to ensure that all aspects of the current design review are conducted satisfactorily
- Solicits input from all participants
- Ensures that all participants maintain proper decorum
- Works with the participants before the meeting to ensure that all required documentation will be available
- Addresses other tasks as necessary to ensure a successful design review

Though it is not mandatory, the DBA typically acts as the leader of design reviews for applications using a database. If the DBA is not the leader, the DBA group must, at a minimum, have the authority to approve the selection of the leader. Other candidates for leader include the application owner or lead developer. Sometimes it is a good idea to hire a consultant who has been

exposed to more applications at many different sites to lead a design review. Doing so can result in the identification of design flaws that might not be caught otherwise. In some cases even a somewhat disinterested third party can be used to drive the review process, as long as that person has experience with design review meetings.

The DBA typically acts as the leader of design reviews for applications using a database.

The Scribe

The responsibility of the scribe is to capture all points of discussion during the design review. Although the scribe is not an active participant in the design review, a scribe is always required. Failure to record the review can result in loss of vital information from the meeting.

The scribe must be capable of understanding the technical discussion but need not have a technical position. The scribe could be a member of the development team who has good writing and listening skills. A technically savvy administrative assistant could be another good choice.

The Mediator

The mediator is an optional role, but depending on the project and the dynamics of the design review team, a mediator can be a blessing. The primary role of the mediator is to negotiate settlements when disagreements occur, and given the nature of a design review, disagreements are almost guaranteed. If a disagreement becomes vocal and volatile, the mediator will hear each side of the disagreement and arrive at an equitable decision.

The primary role of the mediator is to negotiate settlements when disagreements occur.

Although a good leader should be able to resolve most disagreements, the leader's authority may be compromised by the resentment that can ensue from an intervention. By deferring the most difficult and sensitive decisions to the mediator, the leader can maintain the confidence of the group and keep the design review from breaking down. A technical management representative is the usually the best choice for mediator.

The Participants

Design review participants consist of the other actors with a stake in the project. The participants will differ from project to project, and from one design review to the next. From a high-level perspective, though, the following are the recommended personnel to engage in the design review:

Design review participants consist of the other actors with a stake in the project.

- Application development personnel assigned to this development effort
- Representatives from other applications that are affected by the new application or program
- Data administration representative
- Database administration representative
- Representative end users
- End user management
- IT management for the new application and possibly other impacted applications
- Online support representatives for transaction processing and message queuing systems
- Web support personnel for Internet-enabled applications
- Operational support representatives
- IT security representative
- IT audit representative (for regulated and sensitive applications and data)
- Technical support and systems programming representatives

It is not necessary for each of these participants to be involved in each and every facet of every design review. A single application should be subjected to multiple design reviews—with each review focusing on a particular aspect of the application. The scope of each design review should be determined prior to the scheduling of the review so that only the appropriate participants are invited.

Knowledge and Skills Required

To be considered for a position on a design review team, candidates should be experienced in database development. It is best to form the design review

team using participants who possess considerable skills and knowledge. The following criteria should be used to guide the formation of the design review team:

- Strong technical skills: technicians, programmers, and DBAs
- Strong communication skills: all participants
- Good interpersonal skills: all participants
- DBMS fundamentals: all participants to the degree required by their positions
- Background in data modeling and database design: in-depth knowledge for the DA and DBA; good knowledge for programmers and other technicians; some level of knowledge for all other participants
- Strong knowledge of SQL: technicians, programmers, and DBAs

Of course, not every team member will have all of these skills.¹ However, make sure that members of the design review team are experts in their field of practice. For example, an IT manager may have limited expertise in SQL, but that should not exclude him from the design review team. The manager will contribute from his field of experience and should be chosen based on his exposure to the project and his skills as a manager.

Members of the design review team should be experts in their field of practice.

Furthermore, you should strive to maintain the same members of the team throughout the ADLC. Since multiple design reviews are necessary, team member consistency will make design reviews easier because knowledge gained during past design reviews will carry over to subsequent design reviews.

Types of Design Reviews

As previously mentioned, it is best to conduct several design reviews over the course of the ADLC. Multiple design reviews are preferable to a single large design review because they will allow errors and design flaws to be caught earlier in the development process while it is still possible and cost-effective to fix the problems. Additionally, it is unreasonable to expect a single design review to be conducted at the end of the development process because too many details will need to be inspected and errors are more likely to fall

through the cracks.

The following are the seven basic design review phases for a database application:

- Conceptual design review
- Logical design review
- Physical design review
- Organizational design review
- SQL and application code review
- Pre-implementation design review
- Post-implementation design review

Remember the rendering of the ADLC shown in [Figure 1.2](#)? [Figure 6.1](#) points out the relative point within the ADLC where each design review should be conducted.

Initiation

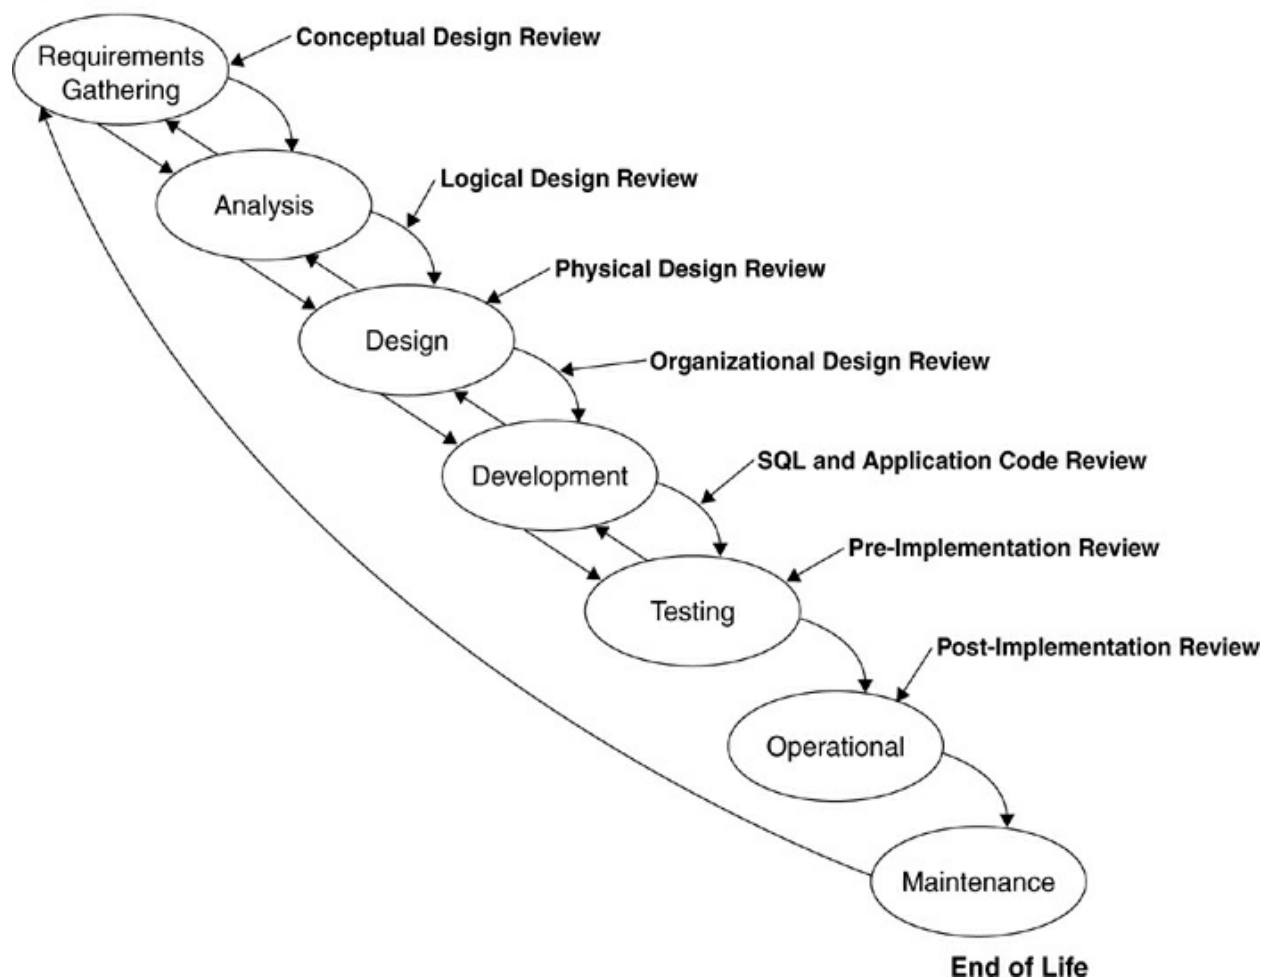


Figure 6.1. Design reviews in the application development life cycle

Conceptual Design Review

The first review to be conducted is the *conceptual design review*. The purpose of this review is to validate the concept of the database and application. The conceptual design review begins with a presentation of an overall statement of purpose and a general overview of the desired functionality to be provided by the application.

The conceptual design review validates the concept of the database and application.

The conceptual design review should be conducted as early as possible in the application development life cycle to determine the overall feasibility of the project. The findings of the conceptual review must verify the purpose of the application and the clarity of the vision for building the databases and supporting application programs.

In order to conduct a conceptual design review a conceptual data model must exist, as well as a high-level design for the application.

Failure to conduct a conceptual design review can result in

- Projects that provide duplicate or inadequate functionality
- Cancellation of projects due to lack of funds, inadequate staffing, poor planning, lack of user participation, or waning management interest
- Projects that run over budget or take longer to complete than anticipated
- Applications that do not deliver the required features and functionality to support the business

The conceptual design review should include the participation of application development, data administration, and database administration staff; end users; and management representatives from the end user team and IT.

Logical Design Review

The *logical design review* follows the conceptual design review. It should be conducted when the first cut of the logical data model has been completed. A thorough review of all data elements, descriptions, and relationships should

occur during this review. The logical design review should address the following questions:

The logical design review examines all data elements, descriptions, and relationships.

- Has the logical data model been thoroughly examined to ensure that all of the required business functionality can be achieved?
- Is the model in (at least) third normal form?
- Have all of the data elements (entities and attributes) required for this application been identified?
- Have the data elements that have been identified been documented accurately?
- Have appropriate data types and accurate lengths been assigned for each attribute?
- Have all of the relationships been defined properly?

The risk of failing to conduct a logical design review is a poorly designed database, which will likely cause data integrity problems. The logical design review helps to ensure that all required data has been identified, designed properly, and fully documented. If changes are made to the logical data model after conducting a logical design review, additional logical design reviews should be scheduled as the project progresses.

Participants in the logical design review should be the same as those who participated in the conceptual design review. If at all possible, the exact same individuals should attend, to maintain a level of consistency from review to review. With the same participants, less up-front preparation will be required because everyone will already be knowledgeable about the purpose of the application and its high-level conceptual design.

Physical Design Review

The *physical design review* comes next—it's the review most often associated with the design review process. In the physical design review, the database is reviewed in detail to ensure that all of the proper database parameter settings and other physical design choices have been made. In addition, the DA and DBA should ensure that a proper translation from logical model to physical database has been made and that all

denormalization decisions are formally documented.

The physical design review ensures that all of the proper database parameter settings and other physical design choices have been made.

The overall operating environment for the application should be described and verified at this stage. The choice of transaction processor and a complete description of the online environment should be provided and verified. An estimation of workload, throughput, and number of concurrent users should be provided and reviewed to ensure that the anticipated requirements can be satisfied. Batch workload should also be reviewed; therefore, a complete description of any batch processes must be provided.

The physical design review may be conducted before all of the SQL that will be used for the application is available. However, general descriptions of all the processes are required to verify the proposed physical database design. Using the process descriptions, the database definitions can be fine-tuned. Furthermore, an initial estimate of whether denormalization could be helpful should be attempted at this point.

Portions of the physical database design may need to be reviewed again as the application development process progresses. Ensuring a valid physical design requires a lot of in-depth attention. Therefore, the review can be broken into discrete processes that can be repeated as changes are made to the database and the application. For example, as SQL statements are written, indexing requirements will change. As indexes are added, the decision-making process should be reviewed to ensure that the indexes are viable for the entire application, not just for a single SQL statement.

Participants in the physical design review should include application development staff, data administration staff, database administration staff, online support representatives, and technical support personnel. If the application or database will affect other applications, or be used by other applications, it would be wise to include representatives from those areas as well.

Organizational Design Review

Smaller in scope than the physical design review, but no less critical, is the *organizational design review*. This review examines the enterprise-wide

concerns of the organization with respect to the new application. The following are some common organizational design review questions:

The organizational design review gauges the impact of the application on the organization.

- How does this system interact with other systems in the organization?
- Has the logical data model for this application been integrated with the enterprise data model (if one exists)?
- To what extent can this application share the data of other applications?
- To what extent can other applications share this application's data?
- How will this application integrate with the current production environment in terms of the DBMS resources required?
- Will the implementation of this application cause the batch window to be exceeded?
- Are the requirements of the application such that online response time and/or data availability are negatively impacted for other users?
- Will the implementation of this application cause the data-processing needs of the shop to expand? For example, will more memory, CPU power, or storage be required?
- Can the organization satisfactorily implement the appropriate level of audit and security required for the application based on industry and government regulations?

Because the purpose of the organizational design review is to gauge the impact of the application on the organization, all the players listed in the "[Participants](#)" section earlier should attend this design review. Failure to include everyone could result in missing certain aspects of the application's impact on the organization because of ignorance or oversight.

SQL and Application Code Design Review

The *SQL design review* is a rigorous review of every SQL statement in the application. Each SQL statement must be reviewed for performance prior to the turnover of the application to production. The review must analyze each statement's access path, the indexes it uses, and possible alternate formulations—resulting in an overall assessment of how it is likely to perform.

The SQL design review is a rigorous review of every SQL statement in the application.

Every DBMS provides a command to show the access path that will be used for a SQL statement. Typically the command is called either EXPLAIN or SHOWPLAN, but I will use EXPLAIN as a generic term. Prior to the SQL design review, an EXPLAIN should be run for each SQL statement. It is important that the EXPLAIN command have access to production statistics. The results of the EXPLAIN statement should be analyzed to determine if the most efficient access paths have been chosen.

Furthermore, every program should be reviewed to validate that efficient programming language constructs were used. Although SQL is more likely to be the cause of poor relational performance, it is quite possible to code an inefficient program using COBOL, Visual Basic, C, or whatever language. For example, a very efficiently tuned SQL statement embedded in a loop within a C program might become very inefficient if that loop runs hundreds or thousands of times. Additional application and SQL performance issues are discussed in detail in [Chapter 12](#), “[Application Performance](#).”

Once again: Every line of code and SQL statement must be reviewed prior to implementation. The SQL and application design review is the appropriate venue for making suggestions for performance improvements prior to moving the application to production status. Alternate formulations and indexing strategies can be suggested during this review and then tested to determine their impact. If better performance is achieved, the application code, SQL, or database design should be modified.

The application developers and DBA are mandatory participants in the application and SQL design review. Additional participants might include application development managers and perhaps technically savvy end users. In some cases, developers may feel more comfortable while their code is being reviewed if their managers are invited. Such an invitation can make it seem less like the DBA is picking apart months and months of the programmer’s hard work. Of course, this decision should be made on a case-by-case basis depending on the developers’ comfort level with their management, the personal interaction skills of the DBA, and the skill level of both the DBA and the programmers.

Pre-Implementation Design Review

A *pre-implementation design review* should be conducted immediately prior to turning over the application to production status. This review consists of an overall appraisal of the system components prior to implementation. Each participant must be prepared to discuss the status of any changes required to support the application once it moves to production. Loose ends existing from previous design reviews should be reviewed to verify that necessary modifications were made and tested. A quick, final review of each application component should be performed to make sure that new problems were not introduced as changes were made.

The pre-implementation design review is an overall appraisal of the system components.

Participants in the pre-implementation design review should include personnel from the application development staff, application development management representatives, database administration staff, online support representatives, and technical support personnel. Optionally, you may choose to invite technically savvy end users.

Post-Implementation Design Review

Finally, we come to the *post-implementation design review*. It is necessary to formally review the application once it has run in the production environment for a while to determine if the application is meeting its objectives, both in performance and in functionality. If any objective is not being met, a plan for addressing the deficiency must be proposed and acted on. Although daily performance monitoring of a new application is a must, it does not preclude the need for a formal post-implementation design review for all new application projects.

A post-implementation design review determines if the application is meeting its objectives.

Because any portion of the application may be a target for improvement, all the players listed in the “[Participants](#)” section of this chapter may be required to attend the post-implementation design review.

Design Review Output

Output from reviews should be clear and concise so that any required

application, SQL, or database modifications can be made quickly and correctly. It is imperative that the scribe capture notes in sufficient detail that a non-attendee can make sense of the discussion. The scribe should edit the notes for grammar and spelling and distribute a copy to all attendees (preferably by e-mail).

Output from design reviews should be clear and concise so that any required modifications can be made quickly and correctly.

An additional result of each design review is a separate list of action items. This list should contain every modification or change discussed during the design review. Each action item should be given a deadline and be assigned to a single person, giving that person the responsibility to make the change, test its impact, and report the progress back to the entire group.

Additional Considerations

There are additional considerations and issues that you will need to deal with as you prepare and conduct your database design reviews. You must be prepared to adapt to changing situations and personnel, as well as to the needs of your organization.

In this section two additional design review considerations are addressed. The first one is how to overcome a potential pitfall: working with a geographically dispersed staff. The second issue is more of an opportunity: using design reviews to mentor junior staff members.

Dealing with Remote Staff

In some cases organizations have distributed workforces where DBAs and development staff are not located at the same site. When staff members are remote, the design review process becomes an even more critical piece of the development project because it forces communication between resources that are not able to interact on a daily basis.

Of course, a distributed staff also complicates the design review. Although it is possible to fly team members to a single location to participate in design reviews, it is rarely cost-effective. Instead of bringing every participant into the same room to conduct the review, a conference call, videoconference, or Web-enabled meeting (such as Live Meeting or WebEx) can be set up. In

such cases it is important that materials be available well in advance of the meeting so that each participant can review the content beforehand.

At any rate, you cannot assume that every staff member who should attend the design review meeting will be located at the same site.

Mentorship and Knowledge Transfer

Design reviews meetings can be a great opportunity for mentoring junior staff members. The meetings should be attended by senior technicians, many with teaching abilities. By inviting junior technicians (who are receptive to new ideas) to the meeting, you potentially can transfer knowledge cost-effectively and efficiently.

Be sure, though, not to turn the meetings into purely education sessions. Furthermore, do not let the junior personnel derail the meeting with endless questions or let the senior personnel use the meeting as a soapbox.

The purpose of the design review is to ensure the viability of the new application and database for the organization. Usually it is possible to do a little mentorship at the same time.

Summary

Design reviews can be time-consuming and difficult to manage, but they are worth the effort. If a systematic approach to database application design reviews is established and followed, the likelihood of implementing optimal applications increases. Database development can be very complex. Only by managing and documenting the implementation process can you ensure the creation of successful and useful application systems. The design review process is an efficient way to encourage a rigorous and systematic pre- and post-implementation review of database applications.

Review

1. Name the roles required for each design review.
2. What are the differences between a logical design review and a physical design review?
3. During which type of design review should denormalization be discussed?
4. Why is it important to review application code in addition to reviewing

SQL?

5. During which phase of the ADLC should the pre-implementation design review be conducted?
6. During which type of design review should the design be checked for conformance to third normal form?
7. Cite several reasons for including representatives from application development management in design reviews.
8. What output is required of every design review?
9. During which type of design review will the impact of the application on the computing resources of the company be ascertained and analyzed?
10. Why should the DBA lead most of the design reviews?

Suggested Reading

DeMarco, Tom, and Timothy Lister. *Peopleware: Productive Projects and Teams*. New York, NY: Dorset House (1987). ISBN 0-932633-05-6.

Freedman, Daniel P., and Gerald M. Weinberg. *Handbook of Walkthroughs, Inspections, and Technical Reviews*. New York, NY: Dorset House (1990). ISBN 0-932633-19-6.

Ginac, Frank P. *Creating High Performance Software Development Teams*. Upper Saddle River, NJ: Prentice Hall (2000). ISBN 0-13-085083-7.

Rothstein, Michael F., and Burt Rosner. *The Professional's Guide to Database Systems Project Management*. New York, NY: John Wiley & Sons (1990). ISBN 0-471-62130-7.

7. Database Change Management

Although a cliché, it is true that change is the only constant in today's complex business environment. An ever-changing market causes businesses to have to continually adapt. Businesses are striving to meet constantly changing customer expectations while trying to sustain revenue growth and profitability at the same time. To keep pace, businesses must constantly update and enhance products and services to meet and exceed the offerings of competitors.

Change is the only constant in today's complex business environment.

Moreover, the individuals within the business usually find it difficult to deal with change. Change usually implies additional roles and responsibilities that almost inevitably make our jobs more difficult. Our comfortable little status quo no longer exists. So, we have to change, too—change either aspects of our environment or our approach to doing things. There are many different elements of managing change, particularly with respect to IT. Each of the following constitutes a different facet of the “change management” experience:

- The physical environment or workplace changes to accommodate more employees, fewer employees, or perhaps just different employees with new and different skill sets.
- The organization changes such that processes or methodology, for example, have to adapt to facilitate a quicker pace for product and service delivery.
- The network infrastructure changes to provide support for a growing, and perhaps geographically dispersed, workforce.
- Applications and systems change to perform different processes with existing data or to include more or different types of data.
- The type and structure of data change, requiring modifications to the underlying database schemata to accommodate the new data.

Change is inevitable but necessary for business survival and success. Many factors conspire to force us into changing our database structures, including

Many factors conspire to force us into changing our database structures.

- Changes to application programs that require additional or modified data elements
- Performance modifications and tweaks to make database applications run faster
- Regulatory changes that mandate storing new types of data, or the same data for longer periods of time
- Changes to business practices, requiring new types of data
- Technological changes that enable databases to store new types of data and more data than ever before

Change will never disappear. Therefore, it is imperative that we have solutions to enable us to better manage these inevitable changes.

Change Management Requirements

To successfully implement effective change management, understanding a set of basic requirements is essential. To ensure success, the following factors need to be incorporated into your change management discipline: proactivity, intelligence, analyses (planning and impact), automation, standardization, reliability, predictability, and quick and efficient delivery.

- *Proactivity*. Proactive change, which can eliminate future problems, is an organization's most valuable type of change. The earlier in the development cycle that required changes are identified and implemented, the lower the overall cost of the change will be.
- *Intelligence*. When implementing a change, every aspect of the change needs to be examined, because it could result in an unanticipated cost to the company. The impact of each change must be examined and incorporated into the change process, because a simple change in one area may cause a complex change in another area. Intelligence in the change management process often requires a thorough analysis that includes an efficient and low-risk implementation plan. True intelligence also requires the development of a contingency plan, should the change or set of changes not perform as projected.
- *Planning analysis*. Planning maximizes the effectiveness of change. A

well-planned change saves time. It is always easier to do it right the first time than to do it again after the first change proves to be less than effective. An effective organization will have a thorough understanding of the impact of each change before allocating resources to implement the change.

A well-planned change saves time.

- *Impact analysis.* Comprehensive impact and risk analyses allow the organization to examine the entire problem, and the risk involved, to determine the best course of action. A single change usually can be accomplished in many different ways. However, the impact of each change may be considerably different. Some changes involve more risks: failure, undue difficulty, need for additional changes, downtime, and so on. All considerations are important when determining the best approach to implementing change.
- *Automation.* With limited resources and a growing workload, automating the change process serves to reduce human error and to eliminate more menial tasks from overburdened staff.
- *Standardization of procedure.* Attrition, job promotions, and job changes require organizations to standardize processes to meet continued productivity levels. An organized and thoroughly documented approach to completing a task reduces the learning curve, as well as the training time.
- *Reliable and predictable process.* When creating any deliverable, a business needs to know that none of the invested effort is wasted. Because time is valuable, a high level of predictability will help to ensure continued success and profitability. Reliability and predictability are key factors in producing a consistently high-quality product.
- *Availability.* Most changes require downtime to implement the change. Applications must come down—the same is true of databases. However, high availability is required of most applications these days, especially for an e-business. This is fast becoming a requirement in the Internet age. Reducing the amount of downtime required to make a change will increase application availability.
- *Quick and efficient delivery.* Consumers demand quick turnaround for most products and services. Profitability is at its best when a product is

first to market. Conversely, the cost of slow or inefficient delivery of products can be enormous. So, when implementing change, faster is better. The shorter the duration of an outage to accomplish the change, the quicker the system can be brought to market.

The Change Management Perspective of the DBA

The DBA is the custodian of database changes. However, the DBA is not usually the one to request a change; a programmer, application owner, or business user typically does that. There are times, though, when the DBA will request changes, for example, to address performance issues or to utilize new features or technologies. At any rate, regardless of who requests the change, the DBA is charged with carrying out the database changes and ensuring that each change is performed successfully and with no impact on the rest of the database.

The DBA is the custodian of database changes.

To effectively make database changes, the DBA needs to consider each of the items discussed in the previous section: proactivity, intelligence, analyses (planning and impact), automation, standardization, reliability, predictability, and quick and efficient delivery. Without a robust, time-tested process that is designed to effect database changes, the DBA will encounter a very difficult job. Why?

Well, today's major DBMS products do not support fast and efficient database structure changes. Each DBMS provides differing levels of support for making changes to its databases, but none easily supports every type of change that might be required. One quick example: Most DBMSs today do not enable a column to be added easily to the middle of an existing row. To accomplish such a task, the DBA must drop the table and recreate it with the new column in the middle. But what about the data? When the table is dropped, the data is deleted unless the DBA was wise enough to first unload the data. But what about the indexes on the table? Well, they too are dropped when the table is dropped, so unless the DBA knows this and recreates the indexes too, performance will suffer. The same is true for database security: When the table is dropped, all security for the table is also dropped. And this is but one example of a simple change that becomes difficult to implement and manage.

Adding to this dilemma is the fact that most organizations have at least two, and sometime more, copies of each database. At the very least, a test and a production version will exist. But there may be multiple testing environments—for example, to support simultaneous development, quality assurance, unit testing, and integration testing. Each database change will need to be made to each of these copies, as well as, eventually, to the production copy. Furthermore, most organizations have multiple DBMS products, each with varying levels of support for making changes. So, you can see how database change can quickly monopolize a DBA's time.

Types of Changes

Managing change is a big component of the DBA's job. In fact, if systems and databases could be installed into an environment that never changed, most of the DBA's job would vanish. However, things change. Business changes usually necessitate a change to application code or to database structure. Less obvious business changes also impact the database—for example, when the business grows and more users are added, when additional data is stored, or when transaction volume grows. Additionally, technological changes such as upgrades to the DBMS and changes to hardware components impact the functionality of database software and therefore require DBA involvement.

Business changes usually necessitate a change to application code or to database structure.

DBMS Software

As discussed in [Chapter 2](#), the DBA must be prepared to manage the migration to new DBMS versions and releases. The complexity involved in moving from one version of a DBMS to another depends on the new features and functions supported by the new version. Additional complexity will be introduced if features are removed from the DBMS in a later version, because databases and programs may need to change if the removed features were being used. Furthermore, as functionality is added to, and removed from, the DBMS, the DBA must create the policies and procedures for the proper use of each new DBMS feature. This aspect of managing change is a significant component of the DBA's job, as we discussed in depth in [Chapter 2](#).

Hardware Configuration

The DBMS may require hardware upgrades or configuration changes. The DBA will be expected to work in conjunction with the system programmers and administrators responsible for setting up and maintaining the hardware. At times the DBMS may require a different configuration from the one that is commonly used, thereby requiring the DBA to communicate to the SA the reason why a nonstandard configuration is required.

Conversely, when hardware changes for other reasons, the DBMS configuration may have to change. Perhaps your organization is changing the disk drives in use with your database server hardware, or maybe adding additional memory to the box. Hardware changes such as these may require changes to database structures and the DBMS configuration. The DBA must be actively engaged with the SA team that configures and maintains the hardware used by the DBMS and, as discussed earlier, may even have to function as an SA in addition to carrying out DBA duties.

Logical and Physical Design

When the database changes, it is important that the blueprints that define the database also change. This means that you need to keep the conceptual and logical data models synchronized with the physical database. This can be accomplished in several ways.

When the database changes, the blueprints that define the database must also change.

Organizations adept in data administration may choose to make changes at the conceptual and logical levels first, and then migrate the changes into the physical database. Usually such an approach requires data modeling tools that segregate the logical and physical models. Furthermore, the procedures must facilitate the specification of changes at each level while providing the capability to synchronize the different models—both forward from logical to physical, and backward from physical to logical.

In the absence of robust data modeling tools, the typical approach taken to synchronize models is manual. Whenever the physical database is modified,¹ the DBA must manually update the logical data model (and perhaps the conceptual data model). Such an effort requires a strict approach to change propagation. As with any manual change management scenario, this approach

is tedious and error prone. However, it is imperative that the logical model be synchronized with the physical database. Failure to do so invalidates the usefulness of the data model as a blueprint for database development.

Applications

Application changes need to be synchronized with database changes; however, this is easier said than done. Whenever changes are made to a physical database structure, application changes usually accompany those changes. For example, simply adding a column to a database table requires application software to populate, modify, and report on the data in the new column.

Application changes need to be synchronized with database changes.

When the database change is migrated to a production environment, the application change must be migrated as well. Failure to do so will render the database change ineffective. Of course, the DBA could allow the database change to be migrated before the application change to ensure that the database structures are correctly specified. After changing the database in the production environment, the DBA would then inspect the database for accuracy and only then allow the application changes to be migrated.

But the relationship between the database change and the application change is valid. If the application change is backed off, the database change should be backed off as well—and vice versa. Failure to synchronize changes will likely cause an application error or inefficiency. It is imperative that the DBA understand these relationships and monitor the change management process to ensure that the database and application changes happen in step with one another.

Physical Database Structures

The most complicated and time-consuming type of change for DBAs is planning, analyzing, and implementing changes to physical database structures. But most databases change over time—indeed, the database that remains static once implemented is very rare. So DBAs must be prepared to make changes to the databases under their care. Some changes are simple to implement, but others are very complex, error prone, and time-consuming. The remainder of this chapter discusses how physical database objects can be

changed and the problems that DBAs can expect to encounter in the process.

The most complicated type of change for DBAs is making changes to physical database structures.

Impact of Change on Database Structures

When the data requirements of your organization change, the databases used to store the data must also change. If the data is not reliable and available, the system does not serve the business—rather, it threatens the health of the business. So, we need infallible techniques to manage database changes. But even more, we need techniques that are not just fail-safe but also automated, efficient, and easy to use. Unfortunately, today’s database systems do not make managing database change particularly easy.

Relational databases are created using Data Definition Language (DDL) statements. DDL consists of three SQL verbs: CREATE, DROP, and ALTER. The CREATE statement is used to create a database object initially, and the DROP statement is used to remove a database object from the system. The ALTER statement is used to make changes to database objects.

Not every aspect of a database object can be changed by using the ALTER statement. Some types of changes require the database object to be dropped and recreated with the new parameters. The exact specifications for what can, and cannot, be changed using ALTER differ from DBMS to DBMS.

For example, you can add columns to an existing table using the ALTER statement, but usually only at the end of the table. In other words, you cannot use ALTER to add a column between two existing columns. Additionally, typically you cannot remove columns from a table. To add a column anywhere but at the end of the column list, or to remove a column from a table, you must first drop the table and then recreate it with the desired changes. Every DBMS has limitations on what can be changed by using the ALTER statement. Furthermore, not just tables but most database objects have certain aspects that cannot be changed using ALTER.

When making changes to a database requires an object to be dropped and recreated, the DBA must cope with the cascading DROP effect. A *cascading DROP* refers to the effect that occurs when a higher-level database object is dropped: All lower-level database objects are also dropped. (See [Figure 7.1](#) for a depiction of the database object hierarchy.) Thus, if you drop a

database, all objects defined in that database are also dropped. The cascading DROP effect complicates the job of changing a database schema.

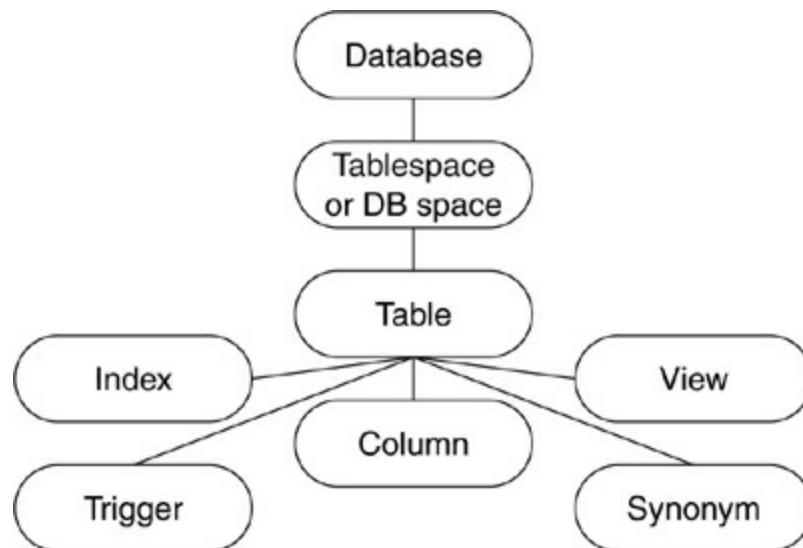


Figure 7.1. Database object hierarchy

The cascading DROP effect complicates the job of changing a database schema.

To understand the complexity involved, let's use an example. Suppose you are working with DB2, and you need to change a segmented tablespace to a partitioned tablespace. To do so, you must drop the segmented tablespace and recreate it as a partitioned tablespace. However, when you drop the segmented tablespace, you are also dropping any tables defined in the tablespace, as well as the table's columns and keys, all indexes defined on those tables, any triggers defined on the tables, any synonyms defined for the table, and all views that access the table. Furthermore, when a database object is dropped, the security information and database statistics are deleted from the system. The DBA must be able to capture all this information prior to dropping the tablespace so that it can be recreated after the partitioned tablespace is implemented. Capturing the DDL from the system catalog or dictionary and ensuring that the DDL is submitted correctly after the modification can be a tedious, complex, and error-prone process.

The system catalog or data dictionary stores the metadata about each database object. Metadata includes much more than just the physical characteristics of database objects. Additional information about database objects such as security authorizations and database statistics are stored along

with the metadata. All of the information required to recreate any database object is readily available if you know where to look for it, and part of the DBA's job is knowing where to look.

A final concern regarding database change: What happens to application program specifications? When database objects that are accessed by a program are dropped, the DBMS may invalidate that program. Depending on the DBMS and the type of program, additional steps may be required to rebind the application to the DBMS after the database objects accessed by the program have been recreated.

The Limitations of ALTER

Many types of database object alteration cannot be performed using the basic SQL ALTER statement; as usual, this varies from DBMS to DBMS and, indeed, from version to version of a single DBMS. However, the actions that are most likely to not be supported by ALTER include

- Changing the name of a database object (depending upon the DBMS and version, some objects can be renamed using the RENAME statement)
- Moving a database object to another database
- Changing the number of tablespace partitions or data files
- Removing a partition from a partitioned tablespace or index
- Moving a table from one tablespace to another
- Rearranging the order of columns in a table
- Changing a column's data type and length (especially to a smaller length)
- Removing columns from a table
- Changing the definition of a primary key or a foreign key
- Adding a column to a table that cannot be null
- Adding or removing columns from a view
- Changing the SELECT statement on which the view is based
- Changing the columns of an index
- Changing whether an index is unique
- Changing whether an index is clustering

- Changing whether the index is ascending or descending
- Modifying the contents of a trigger
- Changing a hash key

In some limited cases, it is possible to use ALTER to change the length of certain types of columns. For example, in DB2 and Oracle you can alter a character column to a larger size, but not to a smaller size. Additionally, it may be possible to change a column from one numeric data type to another. DB2 allows the modification of a column's data type, as long as the change is within the same data type family (numeric to numeric, character to character, or datetime to datetime). For example, it is legal to change a column from SMALLINT to INTEGER using ALTER, but not from SMALLINT to DATE. In general, though, significant changes to the data type and length of a column usually require the table to be dropped and recreated with the new data type and length.

Making physical changes to actual database objects is merely one aspect of database change. Myriad tasks require the DBA to modify and migrate database structures. One daunting challenge is to keep test databases synchronized and available for application program testing. The DBA must develop robust procedures for creating new test environments by duplicating a master testing structure. Furthermore, the DBA may need to create scripts to set up the database in a specific way before each test run. Once the scripts are created, they can be turned over to the application developers to run as needed.

Making physical changes to actual database objects is merely one aspect of database change.

Another challenge is recovery from a database change that was improperly specified, or backing off a migration to a prior point in time. These tasks are much more complicated and require knowledge of the database environment both before and after the change or migration.

The preceding discussion justifies the purchase of a database change management tool to streamline and automate database change management. Keep in mind that the preceding list of items is not exhaustive and that it will differ from DBMS to DBMS. DBA tools exist that manage the change process and enable the DBA to simply point and click to *specify* a change.

The tool then handles all of the details of *how* to make the change. Such a tool removes from the shoulders of the DBA the burden of ensuring that a change to a database object does not cause other implicit changes. Database change management tools provide

- A reduction in the amount of time required to specify what needs to change
- A simpler and more elegant method of analyzing the impact of database changes
- A reduction in technical knowledge needed to create, alter, and drop database objects
- Ability to track all changes over time
- An increase in application availability by reducing the time it takes to perform changes

A database change management tool is one of the first tools acquired by most organizations when they implement a database of any size. Such tools reduce the amount of time, effort, and human error involved in managing database changes. The increase in speed and accuracy when using a change management tool provides an immediate return on the investment to the organization. More information on such DBA tools is provided in [Chapter 23](#), “[DBA Tools](#).”

Database Change Scenarios

A DBA will need to make many different types of changes to a database over its lifetime. Some will be simple and easy to implement, others much more difficult and complex.

As discussed earlier, the SQL ALTER statement can be used to make many types of changes to databases. However, other types of changes may require additional steps to implement. It is the DBA’s job to understand the best way to effect any type of database change. Keep in mind that simple changes often become more difficult in the real world. For example, a simple database change is not quite so simple when it needs to be propagated to multiple databases on different servers at multiple locations.

A simple change is not quite so simple when it has to be propagated to multiple databases on different servers at multiple locations.

A single complex change, such as removing or renaming a column, can take hours to implement manually. Changing the name of one column can require hundreds of changes to be scheduled, executed, and verified from development to test to production. Tackling such challenges is the job of the DBA.

Some Database Change Examples

Adding a new column to the end of a table is usually a very simple type of change. All that is required to implement the change is an ALTER statement, such as

Adding a new column to the end of a table is usually a simple type of change.

[Click here to view code image](#)

```
ALTER TABLE Table_1
  ADD COLUMN new_column INTEGER NULL
;
```

The change can be accomplished in a straightforward manner by issuing a single SQL statement. It simply adds a new integer column to Table_1 that can be set to null. However, making the change once is easy, but keeping track of the change is a little more complex. Tracking database changes becomes more difficult as the number of database environments increases and the latency required between changes increases. In other words, a simple change that needs to be populated across 20 distinct database environments over a period of three months becomes more complex because the DBA must be able to track which environment has which changes. Furthermore, there will usually be multiple changes that need to be tracked.

Tracking database changes becomes more difficult as the number of database environments increases and the latency required between changes increases.

A somewhat more difficult change is modifying the amount of free space for a database object. Such a change typically is accomplished by using an ALTER statement, but additional work is required after the ALTER statement has been issued. For example, consider the following ALTER statement:

```
ALTER TABLESPACE TS1  
    PCTFREE 25  
;
```

This statement changes the free space percentage for the tablespace named TS1 to 25 percent (from whatever value it was before). However, the additional free space does not magically appear after this ALTER statement is issued. In order to reclaim the free space for the tablespace, the DBA will have to reorganize the tablespace after successfully issuing the ALTER statement. Additional work is also required to ensure that sufficient disk space is available for the increased amount of free space. Therefore, the DBA needs to understand how each parameter that can be altered is actually impacted by the ALTER statement. Furthermore, the DBA needs to understand when additional work is required to fully implement the desired change.

Finally, let's examine a very difficult database change: adding a column to the middle of a table. Implementing such a change requires a lot of forethought and planning because it cannot be achieved using a simple ALTER statement. Instead, the table must be dropped and recreated with the new column in the appropriate place. The following steps need to be performed:

Adding a column to the middle of a table is a very difficult change.

1. Retrieve the current definition of the table by querying the system catalog or data dictionary.
2. Retrieve the current definition of any views that specify the table by querying the system catalog or data dictionary.
3. Retrieve the current definition of any indexes defined on the table by querying the system catalog or data dictionary.
4. Retrieve the current definition of any triggers defined on the table by querying the system catalog or data dictionary.
5. Capture all referential constraints for the table and its related tables and determine what their impact will be if the table is dropped (causing all data in the table to be deleted).
6. Retrieve all security authorizations that have been granted for the table

by querying the system catalog or data dictionary.

7. Obtain a list of all programs that access the table by using the system catalog, data dictionary, and any other program documentation at your disposal.
8. Unload the data in the table.
9. Drop the table, which in turn drops any views and indexes associated with the table as well as invalidates any SQL statements against that table in any application programs.
10. Recreate the table with the new column by using the definition obtained from the system catalog.
11. Reload the table, using the unloaded data from step 8.
12. Recreate any referential constraints that may have been dropped.
13. Recreate any triggers, views, and indexes² for the table.
14. Recreate the security authorizations captured in step 6.
15. Examine each application program to determine whether changes are required for it to continue functioning appropriately.

As you can plainly see, such a complex change requires diligent attention to detail to ensure that it is made correctly. The process is fraught with potential for human error and is very time-consuming. In summary, to effectively enact database changes, DBAs must understand all of the intricate relationships among the databases they manage and have a firm understanding of the types of changes supported by the DBMS products they use.

Comparing Database Structures

When managing multiple database environments, the DBA may need to compare one environment to another. Usually changes are made to one database environment, say, the test environment, as applications are built and tested. After the changes have been sufficiently tested, they will be promoted to the next environment, perhaps QA, for additional quality assurance testing. In order to appropriately migrate the required changes, the DBA must be able to identify all of the changes that were applied in the test environment.

One approach to change migration is for the DBA to keep records of each change and then duplicate the changes one by one in the new database environment. However, such an approach is likely to be inefficient. The DBA

could analyze a series of changes and condense them into a single change or perhaps a smaller group of changes, but once again, this approach is time-consuming and error prone.

An alternative approach is to use a DBA tool to compare database components. All of the differences between the environments can be written to a report, or the tool can automatically replicate the structure of the database environment of record to another database environment. To accomplish this, the tool can compare the physical databases using the system catalog, data dictionary, or DDL scripts. A comparison tool is almost a requirement for a very complex database implementation because it is very difficult to track changes from one environment to the next. And the more environments that exist, the more difficult the change management becomes.

Use a DBA tool to compare database components.

If your organization does not have a database change management tool, be sure to save the DDL scripts used to create databases and keep them up-to-date. Every change made to the database must also be made to the DDL scripts. Bear in mind that subsequent ALTER statements can change the database but will not change the DDL scripts. The DBA will need to update the DDL scripts either by appending the ALTER statements to the appropriate DDL scripts or by changing the DDL scripts to reflect the effect of the ALTER statement. Neither approach is ideal: For the first approach, changes may be required that cannot be implemented using ALTER (requiring you to modify the DDL script), and for the second, the likelihood of introducing errors is high because a single change is made twice—once to the actual database and once to the saved DDL script.

If you do not store the DDL scripts for your database objects, you will need to learn how to query the system catalog or data dictionary tables to recreate the database DDL manually. Both of these approaches, saving DDL and manually recreating DDL, are error prone and time-consuming.

Without some type of comparison functionality the DBA must keep track of every single change and accurately record which environments have been changed and which have not. This too is an error-prone process. If the DBA does not keep accurate records, he will have to tediously examine the entire database structure for each database that may have changed, using the system catalog or data dictionary in each database environment. Once again, this is

also an error-prone and time-consuming process.

Requesting Database Changes

The application development team generally requests changes to databases. The DBA is the custodian of the database but is not the primary user of the database. Business users who access data by means of application programs and systems tend to be the primary users of databases.

In order to properly coordinate database changes, the DBA group must institute policies and procedures governing how changes are to be requested and implemented. It is not reasonable to expect database changes to be implemented immediately, or even the same day. However, the DBA group should be held accountable to reasonable deadlines for implementing database changes. The DBA must examine each request to determine its impact on the database and on the applications accessing that database. Only after this information is evaluated can the database change be implemented.

Institute policies governing how changes are to be requested and implemented.

An application developer will request database changes only when those changes are viewed as required. In other words, the application has new data usage needs, and the database needs to be changed to support those needs. Of course, not every request will be implemented exactly as it is received. The DBA may need to modify requests based on his knowledge of the DBMS. Any deviations from the request must be discussed with the development team to ensure that they are still workable within the realm of the application.

Standardized Change Requests

The DBA group should establish standardized forms for implementing database changes. These forms should be customized for each shop, taking into account things such as environment, development expectations, knowledge, DBA experience, production workload, service-level agreements (SLAs), platforms, DBMSs, and naming conventions.

Establish standardized forms for implementing database changes.

Standardized change request forms prevent miscommunications from

occurring during the change management process and, if possible, should be implemented online. The form should include all pertinent information related to each change, including, at a minimum, operating system, database subsystem or instance name, object owner, object name, object type, desired change, and date requested. The form should include sign-off boxes for those personnel who are required to sign off on the change before it can be implemented. Required sign-offs should include at least the application development team leader and a senior DBA—but could also include a business unit representative, DA, or SA, depending on the nature of the request.

When the database change is completed, the form should be signed off by the DBA implementing the change and then sent back to the originator. The originator returns the form to the DBA with a requested date for implementing the change in production.

Checking the Checklists

Many DBAs develop checklists that they follow for each type of database change. These checklists may be incorporated into an online change request system so that the DBA can walk through changes as required. Additionally, many application development teams use checklists to ensure that every step that is required for an application to run correctly is taken and verified.

It is a good practice for the DBA group and the application development teams to share their checklists with each other to verify that each step in the database change and turnover process is completed successfully. Activities performed by DBAs, developers, and technical support personnel often overlap. Allowing the teams to review each other's checklists promotes a better understanding of what steps other units are performing during the change process. Many steps require intricate interaction and communication among the different teams. A formalized review process can correct errors before problems arise.

Share your database change management checklists with other DBAs and the application development teams.

Communication

DBAs must provide education to development organizations on how to request a database change. They should provide guidance on accessing

change request forms, instructions for completing the forms, and guidelines on service expectations.

Unrealistic service expectations, often the biggest problem, can be avoided through education. For example, if requesters understand that the DBA team could take up to two days to process their requests, they will be able to work that delay into their timeline. Clarifying realistic service expectations on an organizational level prevents the DBA team from being deluged by “change this now” requests. These expectations need to be based on solid ground—performance requirements, availability, 24/7 issues, and so on. The DBA must ensure that reasonable time frames—based on solid requirements—are built into policies and procedures that are readily available to all requesters. Only in this way does change management become a discipline.

Unrealistic service expectations can be avoided through education.

Coordinating Database and Application Changes

Although this chapter has focused primarily on the management of database structural changes, DBAs must also participate in the coordination of program changes. Typically, database changes are effected to enable new program functionality. At the very least, new columns will need to be updated and accessed by new program logic.

If the database changes are enacted before the program changes, the new data will not be accessible and existing programs may fail (if the structure changes are incompatible with the existing structure layout in the program).

If the program changes are enacted before the databases changes, the programs will fail because the database structure will not match the expected structure.

It is therefore important for the turnover procedure from the test environment to the production environment to be coordinated and integrated across the DBA group and the application team. The changes should be enacted together. And if the changes need to be backed out for any reason, both should be backed out in conjunction with each other. Failing to do so will result in application and database failures.

Compliance

Regulatory compliance demands need to be integrated into your database change management practices as well. Your change management methodology or tool set must be capable of logging changes in order to track who made what change to which piece of data when. The ability to attest to the purpose of each change is a significant component of regulatory compliance.

Although regulations typically do not impact all database data, it can be difficult to track changes to only database structures that are under the jurisdiction of regulations. Thus, it probably is wise to track and log all database changes to ensure compliance in the event of an audit.

Some forms of database change require that data be unloaded and reloaded. All of the regulations that apply to the production data itself also apply to the unloaded data. Be sure to properly secure and track any data that is extracted, unloaded, archived, or backed up during database changes.

Indeed, regulatory compliance introduces yet another compelling reason to consider adopting a database change management tool. It is possible to develop a process using change management tools to automatically track all structure changes to the database using the change auditing and comparison features of such tools. Capturing all deltas between successive upgrades/changes provides a record of what changed and when to the auditors. These tools also can impose additional security, ensuring that databases are protected with controls to prevent unauthorized changes. Adoption of a change management tool also can contribute to the adoption of a change management policy whereby all changes—routine, nonroutine, and emergency—use the change management tool and follow the appropriate procedures with the requisite level of logging and rigor.

A more complete discussion of regulatory compliance as it applies to database administration is provided in [Chapter 15, “Regulatory Compliance and Database Administration.”](#)

DBA Scripts and Change Management

The most important aspect of database change management for DBAs is orchestrating and implementing changes to database structures without impacting data integrity and availability. However, there are other aspects of change management with which DBAs must contend.

DBAs develop and use many scripts and programs that also should

undergo the rigor of change management. Changes to scripts and processes should be managed, approved, tested, and rolled out in much the same way that changes to production programs are handled. This is especially the case when the DBA team consists of multiple members, perhaps distributed across many locations. When DBA scripts need to be modified, the code should be checked out from a software change management library. By adopting a program change control facility for DBA scripts, you can eliminate problems inherent in parallel development projects, such as dual modification or overwriting changes.

Summary

Databases are guaranteed to require changes over the course of their lifetime. The DBA is the custodian of the database and is therefore responsible for implementing change in a responsible manner that ensures the structure, integrity, and reliability of the database. The DBA must create and administer a database change management discipline consisting of tools, procedures, and policies to effect database change appropriately and responsibly.

Review

1. The DBA is usually the initiator of database changes: true or false? Why or why not?
2. Why is it difficult to add a column between two existing columns in an existing table?
3. Name the three types of database comparisons that may be required in a changing database environment.
4. What is the impact of dropping a database?
5. Describe an alternative change method that can be used if your organization does not use a database change management tool.
6. Explain what is meant by the term *cascading DROP*.
7. Why must the DBA understand the relationship between a database change and the application changes that are needed to use the changed data?
8. Your organization does not use a tool to implement changes to database objects. You need to insert a column between the third and fourth columns of a table with ten columns. Describe the preparation process

you would employ to make such a change.

9. Why would a DBA want to have the capability to compare one database structure to another quickly and simply?
10. If you have to drop an entire table to effect a database change, what other database structures will also be dropped automatically by the DBMS?

Suggested Reading

Ambler, Scott W., and Pramod J. Sadalage. *Refactoring Databases: Evolutionary Database Design*. Boston, MA: Addison-Wesley (2006). ISBN 0-321-29353-3.

O'Donnell, Glenn, and Carlos Casanova. *The CMDB Imperative: How to Realize the Dream and Avoid the Nightmare*. Upper Saddle River, NJ: Prentice Hall (2009). ISBN 978-0-13-700837-7.

Scalzo, Bert, and Dan Hotka. *TOAD Handbook*. 2nd ed. Indianapolis, IN: SAMS/Developer's Library (2009). ISBN 978-0-321-64910-2.

8. Data Availability

Availability is the holy grail of database administrators. If the data is not available, the applications cannot run. If the applications cannot run, the company loses business. Therefore, DBAs are responsible for doing everything in their power to ensure that databases are kept online and operational. This has been the duty of the DBA since the first days of the database.

However, the need for availability is increasing. The age of the long batch window, where databases can be offline for extended periods to perform nightly processing, is over. Contributing to the trend of continual availability is the adoption of e-business. Coupling businesses to the Internet has dramatically altered the way we do business. It has created expectations for businesses to be more connected, more flexible, and importantly, more available. When you integrate the Web with database management, heightened expectations are placed on DBAs to keep databases up and running more smoothly and for longer periods. When your business is online, it never closes. People expect full functionality on Web sites they visit regardless of the time of day. Remember, the Web is worldwide. It may be three o'clock in the morning in New York, but it is always prime time somewhere in the world. Therefore, an e-business must be available and operational 24 hours a day, seven days a week, 365 days a year (366 for leap years). And if your customer is conducting business at three o'clock in the morning in New York, you had better be, too—or you risk losing that customer's business.

It is always prime time somewhere in the world.

On the Web, all of your competitors are just a simple mouse click away. Studies have shown that if Web customers do not get the service they want in seconds, they will take their business elsewhere. A Jupiter Research survey found that 33 percent of broadband shoppers are unwilling to wait more than 4 seconds for a Web page to load. Tests conducted by Amazon revealed that every 100ms increase in load time of Amazon.com decreased sales by 1 percent.

And if prospective customers are satisfied after taking their business to

your competitor, chances are they will never come back. So, an e-business site that is down, even for a short period, will result in not just hundreds or thousands of lost hits, but lost business, too. The impact of downtime cannot be measured in the loss of immediate business alone. No, e-business downtime also damages the goodwill and public image of your organization. And once that is lost, it is quite hard to rebuild the trust of your customers and your company's reputation.

Indeed, some pundits use the phrase "Internet time" to describe the rapid rate of change and the rapid development schedules associated with Internet projects. But the DBA can think of Internet time as a simple Boolean equation—there is uptime and there is downtime. During uptime, business is conducted and customers are serviced. During downtime, business is halted and customers are not serviced. So, Internet-age DBAs are sharply focused on maintaining availability.

Of course, e-business is not the only driver for increased availability. Other factors include

E-business is not the only driver for increased availability.

- The "fast food" mentality of customers who demand excellent service and demand it "now!"
- "Airline magazine syndrome"—you know, when your manager reads an article from the in-flight magazine during his latest junket that states how a competitor offers round-the-clock service . . . so your next project has to offer round-the-clock service, too
- The desire to gain a competitive advantage in the marketplace by offering superior services at a time of the customer's choosing
- The need to react to competitors who offer better service to customers because of higher data availability

Defining Availability

Before discussing further the importance of availability, a good definition of availability is needed. After all, we should know what we are talking about. Simply stated, [*availability*](#) is the condition where a given resource can be accessed by its consumers. This means that if a database is available, the users of its data—that is, applications, customers, and business users—can

access it. Any condition that renders the resource inaccessible causes the opposite of availability: unavailability.

Availability is the condition where a given resource can be accessed by its consumers.

Another definition of [availability](#) is the percentage of time that a system can be used for productive work. The required availability of an application will vary from organization to organization, within an organization from system to system, and even from user to user.

Database availability and *database performance* are terms that are often confused with each another, and indeed, there are similarities between the two. The major difference lies in the user's ability to access the database. It is possible to access a database suffering from poor performance, but it is not possible to access a database that is unavailable. So, when does poor performance turn into unavailability? If performance suffers to such a great degree that the users of the database cannot perform their jobs, the database has become, for all intents and purposes, unavailable. Nonetheless, keep in mind that availability and performance *are* different and must be treated by the DBA as separate issues—even though a severe performance problem is a potential availability problem.

Availability and performance are different and must be treated by the DBA as separate issues.

Availability comprises four distinct components, which, in combination, assure that systems are running and business can be conducted:

Availability comprises manageability, recoverability, reliability, and serviceability.

- *Manageability*—the ability to create and maintain an effective environment that delivers service to users
- *Recoverability*—the ability to reestablish service in the event of an error or component failure
- [Reliability](#)—the ability to deliver service at specified levels for a stated period
- *Serviceability*—the ability to determine the existence of problems,

diagnose their cause(s), and repair the problems

All four of these “abilities” impact the overall availability of a system, database, or application.

Increased Availability Requirements

Talk to the DBA group in any major corporation today and you will hear about an atmosphere of controlled chaos. DBAs are scrambling to address a variety of needs, ranging from the design of new applications to keeping business-critical applications operational. All the while, business executives are demanding that DBAs accomplish these tasks with minimal or no downtime. As more businesses demand full-time system availability, and as the cost of downtime increases, the time available for optimizing performance on business-critical systems and software is shrinking.

The time available for optimizing performance on business-critical systems and software is shrinking.

On the other hand, if routine maintenance procedures are ignored, performance suffers. The DBA is forced to perform a delicate balancing act between the mandate for 24/7 availability and the consequences of deferred system maintenance. The stakes are high, and IT is caught between seemingly contradictory objectives.

The Shrinking Maintenance Window

All growing businesses accumulate enormous amounts of data. And enterprise databases continue to expand, increasing in size and complexity. Multiterabyte databases are common, and the largest databases in production exceed a petabyte in size. At the same time, 24/7 system availability is more a requirement than an exception. DBAs need to be increasingly creative to find time to perform routine system maintenance. High-transaction databases need periodic maintenance and reorganization. With constant use, databases become fragmented, data paths become inefficient, and performance degrades. Data must be put back in an orderly sequence; the gaps created by deletions must be erased. Furthermore, performing defragmentation and reorganization usually results in database downtime.

DBAs need to be increasingly creative to find time to perform routine system maintenance.

Decision Support, Business Intelligence, and Analytics

More and more companies are finding new ways to use core business data for decision support. For example, credit card companies maintain a basic body of information that they use to list purchases and prepare monthly statements. This same information can be used to analyze consumer spending patterns and design promotions that target specific demographic groups and, ultimately, individual consumers. This means that core business data must be replicated across multiple database environments and made available in user-friendly formats. Therefore, the availability of operational data can be negatively impacted by the requirements of decision support users, since large amounts of data are not available for update during bulk data unload processing.

Business intelligence (BI) applications support better business decision making through data analysis. BI technologies provide historical, current, and predictive views of business operations. A BI system and a decision support system (DSS) are somewhat synonymous. Common functions of business intelligence technologies are reporting, online analytical processing, analytics, data mining, process mining, complex event processing, business performance management, benchmarking, text mining, and predictive analytics.

The current buzzwords in this arena, though, are *advanced analytics*, implying a business-focused approach that comprises techniques to help build models and simulations to create scenarios, as well as to understand realities and future states. Advanced analytics uses data mining, predictive analytics, applied analytics, statistics, and other approaches in order to allow organizations to improve their business performance.

The Benefits of Analytics: An Example

A large wireless phone service provider was concerned with the number of customers it was losing. Every customer lost cost the company \$53 in monthly revenue. Although the lost revenue looked small on a customer-by-customer basis, with a large customer base the company was losing millions of dollars each month. Using advanced analytics, it was able to develop an attrition model to predict which customers were most likely to terminate their contract. In doing so, the company developed a model to cross-sell, helping it to retain customers by providing products, services,

and other incentives targeted to their profile. This program improved the retention rate and contributed to an overall savings of \$6.7 million.

That is the type of success story common among companies that have deployed advanced analytics to better understand their data.

Traditional business intelligence enables us to understand the here and now, and even some of the why, of a given business situation. Advanced analytics goes deeper into the “why” of the situation and delivers likely outcomes (refer to the example in the sidebar “[The Benefits of Analytics: An Example](#)”). Although advanced analytics cannot infallibly predict the future, it can provide models for judging the likelihood of events. By allowing business managers to be aware of likely outcomes, advanced analytics can help to improve business decision making with an understanding of the effect those decisions may have in the near future.

Advanced analytics cannot infallibly predict the future, but it can provide models for judging the likelihood of events.

Advanced analytical capabilities can be used to drive a wide range of applications, from operational applications such as fraud detection to strategic analysis such as customer segmentation. Regardless of the applications, advanced analytics provides intelligence in the form of predictions, descriptions, scores, and profiles that help businesses better understand customer behavior and business trends. And decision support, business intelligence, and advanced analytics all rely upon having access to data—large amounts of data.

Data Warehousing

Just as decision support, business intelligence, and advanced analytics have expanded the use of operational data, data warehousing has driven overall database growth. Typical data warehouses require the replication of data for use by specific departments or business units. The unloading and loading of external data to operational data stores, and then on to data warehouses and data marts, has increased the number of database utility operations that need to be run and administered. The time taken to propagate data has conversely affected the overall availability window of both the data sources and data targets during unload and load processing. The growth of data warehouses

will continue unfettered into the foreseeable future, fed by the information needs of knowledge workers, the expanding business embrace of advanced analytics, and the falling cost of storage media.

The growth of data warehouses will continue unfettered into the foreseeable future.

Full-Time Availability

Just when the latest hardware and software technologies are beginning to bring 24/7 availability within reach, the mandates of the global economy have forced IT departments to reevaluate the situation. Now the buzz phrase is 24/24 availability, because businesses conduct operations in all time zones and data must be available to a new spectrum of users, not all of whom work in the same time zone as the operational DBMS.

Airline reservation systems, credit card approval functions, telephone company applications—all must be up and running all day, every day. International finance is one of the best examples of the need for full-time availability. Money never sleeps, and the daily flow of euros, dollars, pounds, and yuan occurs daily and inevitably. So does the global information exchange on which brokers base their buy and sell decisions. Large quantities of money are on the line every minute, and downtime simply cannot be tolerated. DBA and IT professionals need techniques that perform maintenance, backup, and recovery in small fractions of the time previously allotted to accomplish these tasks.

Increasing IT Complexity

Any single-vendor system should be clean, precise, and predictable. But today, it is hard to find a company of any size that does not operate in a heterogeneous environment that includes mainframe, midranges, and desktop systems in a client/server infrastructure. As these systems expand in size and functionality, IT staffs must find ways to accommodate operational tuning across a complex, heterogeneous IT environment. This is rarely the seamless process portrayed by the hardware manufacturers. Moreover, the DBMS software itself also can add complexity, with new releases and features being delivered at breakneck speed.

It is hard to find a company of any size that does not operate in a heterogeneous environment.

Complexity stems from human factors as well. Downsizing has forced former IT specialists to become generalists. As a result, tasks such as database reorganization—something that used to be simple and straightforward for expert DBAs—are now complex and lengthy for generalists. Of course, IT is not immune to corporate downsizing; there are now fewer personnel to handle day-to-day computer issues than there were just a few years ago. Finally, mergers and acquisitions force IT staffs to consolidate incompatible systems and data structures.

Cost of Downtime

The cost of downtime varies from company to company. Contingency Planning and Research (a division of Edge Rock Alliance Ltd.) estimates it at approximately \$6.5 million per hour at retail brokerage houses. Consult [Table 8.1](#) for additional examples of the estimated hourly cost of downtime by industry. Of course, these numbers are approximations and estimates—each organization needs to determine the actual cost of downtime based on its customers, systems, and business operations.

Table 8.1. Cost of Downtime

Type of Business	Estimated Hourly Financial Impact of an Outage
Retail brokerage	\$6.45 million
Credit card sales authorizations	\$2.6 million
Home shopping channel	\$113,750
Catalog sales centers	\$90,000
Airline reservation centers	\$89,500
Package shipping service	\$28,250
ATM service fees	\$14,500

Some businesses can handle downtime better than others. For brokerages, downtime is a catastrophe. For other businesses that can “get by” using manual systems during an outage, downtime is not as much of a disaster. The truth is, outages impact every business, and any nontrivial amount of downtime will impose a cost on the organization. When estimating the cost of

downtime, remember to factor in all of the costs, including

Outages impact every business.

- Lost business during the outage
- Cost of catching up after systems are once again available
- Legal costs of any lawsuits
- Impact of reduced stock value (especially for dot-coms that rely on computerized systems for all of their business)

Additionally, downtime can negatively impact a company's image. In this day and age an outage of any length that impacts business, particularly e-business, will be reported by the press—and if the story is big enough, not just the computer press but the business press as well. Let's face it, bad news travels fast. Recovering from negative publicity can be a difficult, if not impossible, task.

Sometimes companies are unwilling to spend money on software and services to improve availability because they do not have an understanding of the true cost of downtime for their business. One line of thought goes something like this: "I know our systems may go down, but the chance of it impacting us is really small, so why should we incur any costs to prevent outages?" Such thinking, however, can be changed when all of the cost and risk factors of downtime are known and understood.

Failure to prepare an estimate of the cost of downtime will make it more difficult to cost-justify the measures a DBA needs to take to ensure data availability.

How Much Availability Is Enough?

Availability is traditionally discussed in terms of the percentage of total time that a service needs to be up. For example, a system with 99 percent availability is up and running 99 percent of the time and down, or unavailable, 1 percent of the time. Another term used to define availability is MTBF, or mean time between failure. More accurately, MTBF is a better descriptor of reliability than availability. However, reliability has a definite impact on availability. In general, the more reliable the system, the more available it will be.

In this Internet age, the push is on to provide never-ending

uptime.

So, just how much availability is enough? In this Internet age, the push is on to provide never-ending uptime, 365 days a year, 24 hours a day. At 60 minutes an hour, that mean 525,600 minutes of uptime a year. Clearly, to achieve 100 percent availability is a laudable goal, but just as clearly an unreasonable one. The term *five nines* is often used to describe highly available systems. Meaning 99.999 percent uptime, five nines describes what is essentially 100 percent availability, but with the understanding that some downtime is unavoidable (see [Table 8.2](#)).

Table 8.2. Availability versus Downtime

Availability	Approximate Downtime per Year	
	In minutes	In hours
99.999%	5 minutes	0.08 hours
99.99%	53 minutes	0.88 hours
99.95%	262 minutes	4.37 hours
99.9%	526 minutes	8.77 hours
99.8%	1,052 minutes	17.5 hours
99.5%	2,628 minutes	43.8 hours
99%	5,256 minutes	87.6 hours
98%	10,512 minutes	175.2 hours (or 7.3 days)

Even though 100 percent availability is not reasonable, some systems are achieving availability approaching five nines. DBAs can take measures to design databases and build systems that are created to achieve high availability. However, just because high availability can be built into a system does not mean that every system should be built with a high-availability design. That is so because a highly available system can cost many times more than a traditional system designed with unavailability built into it. The DBA needs to negotiate with the end users and clearly explain the costs associated with a highly available system.

Whenever high availability is a goal for a new system, database, or application, careful analysis is required to determine how much downtime

users can really tolerate, and what the impact of an outage would be. High availability is an alluring requirement, and end users will typically request as much as they think they can get. As a DBA, your job is to investigate the reality of the requirement.

High availability is an alluring requirement, but your job is to investigate the reality of the requirement.

The amount of availability that should be built into the database environment will be based on cost. How much availability can the application owner afford? That is the ultimate question. Although it may be possible to achieve high availability, it may not be cost-effective, given the nature of the application and the budget available to support it. The DBA needs to be proactive in working with the application owner to make sure the application owner fully understands the cost aspect of availability.

Availability Problems

Since the focus of this book is on database administration, it would be fair to assume that DBAs must manage the availability of the databases under their control. Although this is true, it is an incomplete definition of the DBA's duties with regard to availability. To further understand why this is so, let's first examine all of the potential causes of data unavailability.

Loss of the Data Center

Quite obviously, data will not be available if the data center is lost due to a natural disaster or some other type of catastrophe. Whether the disaster is small from a business perspective and impacts only the computing resources of the business or whether it is large and impacts an entire building or the entire organization, losing the computer means losing the database and any data it contains.

To restore availability in a disaster situation usually requires recreating the database environment (and perhaps much more) at a remote location. Preparing for such a scenario is covered in detail in [Chapter 17, "Disaster Planning."](#)

From an availability perspective, losing the data center is the worst type of availability problem the DBA could ever encounter. Even after the data and databases have been restored at the remote location, serious availability

issues will linger. For example, probably all of the data will not be up-to-date, which may require the DBA and the users to recreate data before allowing general access to the databases. Additionally, ensuring that users have the proper connections to the new location may increase the outage, and problems with an unfamiliar setup can cause subsequent outages. Planning for disasters and developing a good disaster plan will minimize such problems.

Losing the data center is the worst type of availability problem a DBA could encounter.

Network Problems

If a database server is on a single network switch and that switch incurs an outage, the database will be unavailable. Consider implementing redundant network switches to prevent such outages.

Loss of network access also can cause a database outage. Such problems are usually caused by malfunctioning hardware, such as the network card in the database server. It is wise to have spare networking hardware available for immediate replacement in case such a problem occurs.

However, not all network problems are hardware problems. Installing a new version of the networking software or specifying inaccurate network addresses can cause database outages. The DBA is not a networking expert and cannot be expected to resolve all types of network problems; however, the DBA should be able to recognize when a problem is caused by the network rather than the DBMS, the application, or some other piece of hardware or software used by database applications.

DBAs should build good working relationships with the networking specialists in their organization. When networking problems occur, the DBA can consult with the network specialists right away and ideally learn something from the experience of working with them.

Loss of the Server Hardware

At a basic level, the database server hardware consists of the CPU, the memory, and the disk subsystems holding the databases. Let's examine how the loss of any or all of these components impacts database availability.

Obviously, if the CPU is damaged or becomes unavailable for any reason,

the database also will not be available. This is true even if the CPU is the only piece of the database server that is lost. Even if system memory and the disk subsystem remain intact, the database will be inaccessible because the CPU drives all computer processes. However, the database files should remain usable and could possibly be connected to another CPU to bring the database back online.

To avoid outages due to CPU failure, consider using hardware *cluster failover* techniques. When using cluster failover, the loss of a single server causes the system to process on another node of the cluster. The data need not be moved, and failover is automatic.

Another approach is to use a standby system: A copy of the database logs produced on the primary server is shipped to the secondary server, or data from the primary server is replicated to the secondary server. An alternate approach is to keep a second server configured identically to the primary one so that the drives can be pulled out of the primary server and simply inserted into the secondary server.

Any such rapid failover approach will depend on the type of hardware you are using, the ability of the hardware to participate in certain failover techniques, and the cost associated with the technique. Refer to the sidebar [“Oracle Standby Database Options”](#) for one approach to database failover.¹

If system memory fails, the database may or may not be available. If all system memory fails, any database on the system will be unavailable because memory is required for a relational database system to operate. To resolve this situation you will need to replace the failing RAM chips or modules. However, the database may remain available when only some of the system memory fails, although it is likely that performance will suffer. Once again, replacing the faulty memory components should resolve the problem.

If the entire database server is lost or damaged, the failure will be more difficult to address. Loss of an entire server implies that the CPU, memory, and disk subsystem are unavailable. Once again, in such a situation the databases on the server also will be unavailable. If the entire server platform fails, the database will need to be recreated. It is not safe to assume that the data on the disks can be accessed or would be valid even if it could be accessed. The entire database system, including the configuration, connections, and data, would need to be rebuilt. Data would be as accurate and up-to-date as the last database log that is available for the system.

Database files would need to be recovered from backup copies and the changes on the database log reapplied. For more information on database backup and recovery, refer to [Chapter 16](#).

Oracle Standby Database Options

Oracle supports the ability to create standby databases. A *standby database* is a copy of a functioning database that is kept ready for use in case the original database incurs an outage. Typically, the original database is referred to as the primary, or source, database, and the standby database is referred to as the secondary, or target, database.

It works like this: The transactions applied to the primary database are applied to the standby database using the redo logs of the primary database. When Oracle archives the redo logs, the archived logs are sent to the standby database server and applied. Therefore, the standby database will always be a little less current than the primary database.

If the primary database suffers an outage, the DBA can switch to the standby database. This is done by taking the standby database out of recovery mode, applying all archived redo logs, and then, if possible, applying the transactions from the current redo log. The standby database is then switched to become the primary database. When the original primary database is repaired, it becomes the new standby database.

Oracle also supports standby instances and standby tables when different degrees of availability are required. Of course, there are more implementation details and specifics that an Oracle DBA needs to understand before implementing standby databases. Such implementation-specific details are beyond the scope of this book and the reader is directed to the Oracle product documentation for details.

Keep in mind, though, that the implementation of a standby database does not eliminate the need to make backup copies of either the primary or secondary databases. Nevertheless, standby Oracle databases can help to increase the level of availability in the event of unplanned outages.

Losing an entire database server is rare if precautions are taken, but it is insufficient simply to purchase a reliable server and forget about downtime.

Hardware goes bad regardless of built-in ruggedness and availability features. Power spikes, power outages, and human error are the most likely causes of such server failure. The DBA should always ensure that redundant power supplies and uninterruptible power supply (UPS) systems protect the database server against sudden power loss—and that proper precautions are taken such that any type of outage incurred due to server hardware failure is minimized.

Losing an entire database server is rare if precautions are taken.

Disk-Related Outages

Because databases rely on physical disk structures to actually store the data, database availability is particularly vulnerable to disk failures. Of course, the degree of database unavailability depends on the type of disk system and the type of outage the disk system has suffered. Disk drives fail for numerous reasons: The physical drive mechanism may fail, the controller could fail, or perhaps a connecting wire has loosened.

Database availability is particularly vulnerable to disk failures.

One of the simplest database implementations is to store the database files on a local disk subsystem attached directly to the database server. Obviously, if the local disk subsystem fails, the database becomes unavailable. Recovery from such a failure typically requires the server to be reconfigured with a new disk subsystem to which the database is restored. Another recovery method is to bring up an entirely new database server and restore the database there. Neither of these options is required if the failing disk subsystem can be fixed.

To restore the database, the DBA can recover the data using backup copies of the database files and logs. RAID disk systems can help to minimize outages as well because multiple disk drives must fail before an outage occurs. [Chapter 18, “Data Storage and Management,”](#) provides further details on RAID storage.

Another scenario is a database stored on a storage area network (SAN). A SAN is a grouping of networked disk devices. If the SAN fails, the database will likely become unavailable. SAN failure can occur if multiple disk drives within the network fail, the SAN connections fail, or a power loss occurs.

A failover system, such as a standby database, can help to minimize the

outages associated with disk-related failures. With a standby database, the only data likely to be lost because of a disk failure is from uncommitted transactions at the time of the failure.

Operating System Failure

Not all database availability problems are caused by hardware problems. Software also can be the culprit. For example, data will not be available during an operating system (OS) failure or outage, even if all of the server hardware is operational. Typical causes of operating system outages include general OS instability due to inherent bugs, problems encountered when upgrading an OS version, or problems with patches applied to the operating system.

Data will not be available during an OS failure or outage, even if all of the server hardware is operational.

When an OS failure occurs, the only viable options for restoring data availability are to fix the OS problem or to restore the database on another server with a functional operating system. Once again, a failover system can be used to minimize downtime caused by a software failure.

DBMS Software Failure

Similar to the failure of an operating system, a failure in the DBMS software will cause unavailability. If the DBMS is not operational, the data in its databases cannot be accessed. DBMS failure occurs for similar reasons that an OS fails: general DBMS instability due to inherent bugs, problems encountered when upgrading to a new version of the DBMS, or problems when patches are applied to the DBMS software. The DBMS may also fail when resources it needs to operate are not available—such as start-up parameters, certain system files, and memory structures. For example, if the database log file is corrupted or missing, the DBMS will not allow data to be modified.

If the DBMS is not operational, the data in its databases cannot be accessed.

In short, the DBMS will fail when its software has bugs or it cannot gain control of the resources it needs to operate correctly.

Application Problems

Application software also can cause availability problems. If the application is the only way that end users can access data, a data unavailability problem arises when the application fails. An application software failure is unlikely to cause a database outage, but an operational database means little if the users can't use their usual programs to access the data. The DBA and other sophisticated users will be able to access data in the database using alternative methods such as interactive SQL or query tools.

Software bugs and glitches or the loss of program execution modules or libraries can cause application outages. Thorough program testing and quality assurance procedures can minimize the occurrence of application outages.

Security and Authorization Problems

Security-related problems are another cause of database unavailability. This type of problem is caused by improper use or administration of the database. Before data can be accessed, either directly by an end user or by means of a program, authorization must be granted—typically by the DBA or security administrator. If, for any reason, the requesting agent does not have current authorization to access the database, the data will be unavailable to the requester. Security-related problems usually occur immediately after an application goes into production or when a new user attempts to use the system but has yet to receive permission. DBA error can also cause security-related problems: If the DBA accidentally overwrites or removes valid authorizations from the DBMS, valid users will not be able to access data.

Make sure that proper security and authorization procedures are followed at your site.

To avoid security-related problems, make sure that proper security and authorization procedures are followed at your site, and use extra caution when changing database security. More information on this topic is provided in [Chapter 14](#), “[Database Security](#).”

Corruption of Data

Corrupt data can cause unavailability. Even if the hardware and system software are intact, a database can become unavailable because its contents are inaccurate. Business decisions based on faulty or corrupt data can

negatively impact the business. Corrupt data comes from many sources: application program bugs, DBMS software bugs, bad database design, or user error. Old data also can be faulty if newer data was subjected to more rigorous quality and edit checking.

Corrupt data can cause unavailability.

If a sufficient amount of the data becomes corrupt, the DBA will need to take the database offline. Most DBMS systems provide commands to do this. If the database is not taken offline, not only will bad business decisions be made but the data may become more corrupt and processes that use the data may corrupt other, related data. The sooner the DBA takes action to make the data unavailable, the less damage that will ensue. Needless to say, taking a database offline equals lost business.

When data is corrupted, the DBA must work with the application specialists to identify exactly which data is inaccurate and to develop a plan to correct the data. Furthermore, the team must identify the cause of the corruption and take action to correct it. This very time-consuming task requires a significant amount of effort.

After identifying the corrupt data elements, the DBA and application specialists may be able to restore access to the unaffected portion of the database while the bad data is cleansed. However, the database should not be brought back online for update processes until the process that was corrupting the data has been identified and corrected.

A different type of availability problem can be caused when data transmissions fail. If users rely on the data in the database to be accurate and up-to-date, a failed data feed can render the database useless, even if there is accurate data in the database, because it will be impossible to determine which data is up-to-date and which is not. Such a failure is similar to data corruption in that the data exists but is not usable. Correcting such problems, though, is usually many times easier than correcting invalid data. At any rate, the DBA must work with the network technicians and SAs to identify and correct data transmission failures. Once the feed is available, the DBA needs to load the new data into the database quickly to resolve the availability problem.

Loss of Database Objects

Inadvertently dropping a database object is yet another cause of unavailability. When a tablespace or table is dropped, the data is no longer accessible. The DBA will have to run scripts to recreate the objects (including any related objects that were dropped), reapply any referential constraints, rebuild authorizations, and then reload the data, which may physically still reside on the disk device. Some DBMS products and add-on tools facilitate easy recovery of that data. If such techniques are not available, the data will need to be recovered from backup files and archive logs.

Inadvertently dropping a database object is another cause of unavailability.

If an index is dropped, the data will be accessible but performance will usually suffer. If performance is bad enough, the user may experience the equivalent of an unavailable database. For example, consider a process that returns a single row from a multimillion-row table. If an index is used, access is rapid because the index is read using a few I/O operations and then a direct pointer to the table data is used to perform a single I/O (usually) to read the data. Without the index, the DBMS will have to scan all the multimillion rows to return the single requested row—and that can take hours, if not days, depending on the DBMS, hardware, and system environment.

Dropping a view will cause availability problems for those users and processes that rely on the view for access to the database. Dropping a view does not delete any data from the database; it simply removes the view definition from the system. However, if a user accesses the database only by means of the view, removing the view results in database unavailability for that user. To rectify the situation, the DBA will need to recreate the view, reapply the view authorizations, and possibly rebind application packages or plans (depending on the DBMS and how that DBMS is being used).

Dropping database objects mistakenly usually is the result of simple human error on the part of a DBA. However, it may also result when inappropriate or unskilled individuals are granted DBA-like authority on the database. You can minimize inadvertent dropping of database objects by creating and ensuring appropriate database security and providing in-depth training for DBAs. Third-party tools are available that automate the restoration of dropped database objects.

Loss of Data

It is possible for data to be unavailable because it no longer exists, which can occur if data is deleted or overwritten erroneously. Accidental mass deletes, application program bugs, or even malicious attacks can all cause data loss.

Accidental mass deletes, application program bugs, or even malicious attacks can all cause data loss.

When data is deleted in error, the DBA may need to recover the database to a point in time before the data was deleted. This can be accomplished using the RECOVER or RESTORE function of the DBMS in use. The DBA reviews the transaction logs to determine when the error occurred and then restores the data to the point in time just before the error. Analyzing and interpreting the contents of a database transaction log can be a difficult proposition, but there are third-party tools on the market that simplify the task.

If you are using a standby database, there is a slight possibility that you can avoid an outage when data is lost. If there is a sufficient delay in the process that ships the log to the secondary database, the error may not yet have been introduced to the standby database. In that case, you can switch the processing to the secondary database and correct the primary offline.

Data Replication and Propagation Failures

Many databases participate in replication and propagation processes to synchronize data that resides in separate databases. If a replication or propagation task fails, the data will still be accessible, but it may not be as current as the users expect it to be.

There are really two scenarios for replication and propagation failure: one for databases that participate as a subscriber and another for databases that participate as a publisher. Subscribers receive replicated or propagated data from publishers. It is possible that data in the subscriber database will not be up-to-date because the publisher fails to replicate or propagate changes. Such problems can occur for many reasons, including connectivity problems, software failures, or scheduling errors. At any rate, availability problems will result if the data is not up-to-date in the subscriber database.

A second type of problem can occur with the publishing database. If the publishing task fails, the publisher will be up-to-date, but once the replication or propagation is fixed and the service is reinitiated, performance may suffer

as the publisher attempts to catch up by sending large amounts of data to the subscriber. This, too, can cause availability problems if the performance degradation is severe enough. Regardless of the scenario, the DBA needs to develop procedures to detect replication and propagation outages and methods to rapidly restore the replication or propagation process.

Severe Performance Problems

A severe performance problem can cause unavailability. Even if the database is technically available—with accurate data and running on operational hardware—poor performance can make the database unusable. Any number of problems can cause poor performance, including damaged indexes, improperly defined indexes, data growth, additional users, out-of-data database statistics, and locking problems. The end users do not care what the reason is—they just want their applications to work. Therefore, even though a performance problem is not technically the same as an availability problem, the DBA must treat it as such if the problem is severe enough to cripple the end users' ability to access data.

Poor performance can make the database unusable.

For more information on monitoring, tuning, and managing database performance refer to [Chapters 9](#) through [12](#).

Recovery Issues

Your organization's database backup and recovery strategy will have an impact on the availability of your database systems. There are multiple techniques and methods for backing up and recovering databases. Some techniques, while more costly, can enhance availability by recovering data more rapidly.

It is imperative that the DBA team create an appropriate recovery strategy for each database object. For data with high availability requirements, the backup and recovery strategy should provide the shortest mean time to recover. When it is imperative that the database environment be recovered as quickly as possible after a database crash, using the most effective backup and recovery techniques is critical. Factors that impact recoverability and reinstating availability include operating system configuration, hardware architecture design, database features, backup frequency, and recovery procedures and practices.

Many modern disk arrays offer instant copy and recovery of data using dynamic mirroring and other inventive techniques. By using storage hardware technology, a DBA can significantly reduce the impact of database backup recovery. Be sure to investigate the storage solutions available to you and to use their built-in functionality for improving data availability.

Additionally, be sure to investigate the high availability and recovery features of your DBMS when building applications that require extreme availability. For example, refer to the sidebars on DB2 HADR and SQL Server 2012: AlwaysOn.

Refer to [Chapter 16](#) for in-depth coverage of database backup and recovery practices, as well as procedures for improving database availability.

DBA Mistakes

One of the biggest causes of database downtime is human error. In fact, one major DBMS vendor states that 70 percent of the calls it receives for database outages are the result of DBA mistakes. While there is nothing you can do to guarantee that mistakes will not be made, proper DBA training and tools can minimize mistakes.

One of the biggest causes of database downtime is human error.

DB2 HADR

DB2 high-availability disaster recovery (HADR) is a data replication feature that provides a high-availability solution for both partial and complete site failures. It is a feature of DB2 on Linux, UNIX, and Windows platforms that was ported to DB2 from Informix.

HADR protects against data loss by replicating data changes from a source database, called the primary, to a target database, called the standby. Without HADR, the DBMS server has to be rebooted when a failure occurs. The time required to restart the database and the machine where it resides can be highly unpredictable, ranging from a few minutes to much longer. With HADR, a standby database can take over in seconds. Clients can be redirected to the standby database using automatic client reroute or retry logic in the application.

HADR allows the standby database to take over as the primary

database with full DB2 functionality. It is also possible for the original primary database to be brought back up and returned to its status of primary database.

SQL Server 2012: AlwaysOn

SQL Server 2012 delivers improved availability through AlwaysOn SQL Server Failover Cluster Instances and AlwaysOn Availability Groups. These AlwaysOn features deliver improved database uptime and data protection.

AlwaysOn SQL Server Failover Cluster Instances enables multisubnet failover clusters, which are configurations where each failover cluster node is connected to a different subnet or different set of subnets. There is no shared data storage, so data needs to be replicated between the multiple subnets. This helps to improve disaster recovery as well. Flexible failover policy provides for moving group ownership to another node in the case of a failure.

AlwaysOn additionally provides for the creation of Availability Groups, which are containers that define a set of databases to failover as a single unit, as well as a set of availability replicas to host copies of each availability database. AlwaysOn Availability Groups can be deployed to improve database availability through the creation of multiple secondary replicas, choosing between asynchronous-commit and synchronous-commit mode, and deployment of various failover modes, among other benefits.

Be sure that all DBAs have received proper training before giving them responsibility for critical-production database systems. Training should consist of both course material (either instructor-led or computer-based training) and on-the-job experience. A DBA's first administration experience should always be with a test system—never with production data.

Additionally, third-party tools can greatly diminish the effort required to keep databases up and running. Database performance monitors can be used to immediately notify DBAs of outages. Simple problems such as “out of space” conditions can be eliminated by proactive performance tools because they can be set up to warn about a potential problem before the problem

occurs. The DBA can take corrective action when notified, and a potentially devastating problem that would cause an outage can be averted.

Some database performance tools provide proactive triggers that can be set up to automatically correct problems as they occur. Other DBA tools, such as recovery managers and advanced reorganization tools, can be used to fix problems rapidly when they are encountered. More information on the vast array of available DBA tools can be found in [Chapter 23](#).

Outages: Planned and Unplanned

When discussing outages and downtime, many technicians immediately think of unplanned outages. Human error, software bugs and glitches, and hardware failures cause unplanned outages. However, most downtime actually is caused by planned outages (see [Figure 8.1](#)). Planned outages are caused by regularly scheduled system and database maintenance tasks that require the database to be offline—for example, database reorganization or changes.

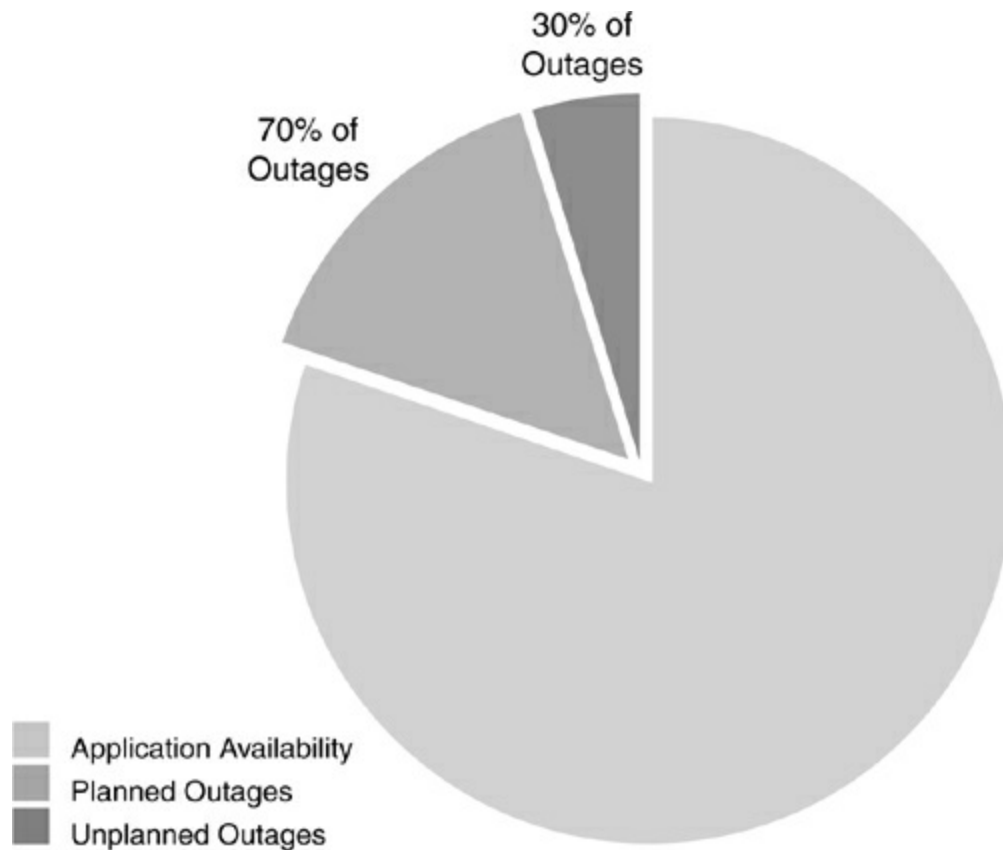


Figure 8.1. Planned and unplanned outages

Most downtime actually is caused by planned outages.

Planned outages represent as much as 70 percent of downtime; unplanned outages represent the other 30 percent. Furthermore, studies show that as much as 50 percent of unplanned downtime is due to problems encountered during planned downtime. Therefore, if the majority of outages are caused by planned actions, it makes sense that DBAs should concentrate more effort on developing techniques to avoid outages during planned database changes and maintenance activities. It's a given that DBAs need to prepare for the inevitability of downtime due to unplanned outages. However, because planned outages are actually a greater risk to availability, the DBA can have a more beneficial impact by developing techniques to reduce them.

Ensuring Availability

Now that we have established that life is challenging for today's DBAs, we will shift our focus to some techniques that help promote higher availability. Faced with shrinking budgets and resources, and an ever-increasing volume of data to manage, IT organizations need to evaluate their critical needs and implement a series of key strategic steps to ensure availability. Good strategy should include steps to

- Perform routine maintenance while systems remain operational
- Automate DBA functions
- Exploit the features of the DBMS that promote availability
- Exploit hardware technologies

Perform Routine Maintenance While Systems Remain Operational

To address the need for performance optimization while trying to get the most out of smaller IT staffs and budgets, products that simplify and automate maintenance functions are key. DBAs need tools that reduce maintenance time from hours to minutes or no time at all, while allowing users continued access to the data they need to do their jobs. Some DBMS products provide built-in features to perform some maintenance tasks while the database is available. If the DBMS does not provide native support, tools are available from ISVs that provide additional database availability. The key is to use the functionality of nondisruptive database utilities.

Use the functionality of nondisruptive database utilities.

A *nondisruptive utility* is a task that provides both update and read access to a database during execution of database maintenance—and does so without a loss of data integrity. Considerations for deployment of nondisruptive utilities are the number and types of resources needed to perform nondisruptive operations. In general, native database utilities use considerably more CPU and I/O resources than ISV utility solutions, but the ISV utilities can be costly. The types of nondisruptive utilities that are needed most include

- Database reorganization, to maintain performance
- Database backup, to ensure data is available for recovery in the event of application or hardware failure, in addition to disaster recovery preparedness
- Database recovery solutions that can apply recovered data without requiring an outage
- Unloading and loading processes for moving data between source data and operational data stores for decision support systems and data warehouses
- Statistics-gathering utilities that can analyze data characteristics and record statistics for use by the database optimizer
- Integrity-checking utilities for both referential integrity and structural data integrity

One example is online database reorganization. Typically, an online REORG is performed by making a duplicate copy of the data and reorganizing the duplicate. Read and write access continues on the original data. When the reorganization of the shadow copy is complete, the REORG process uses the database log to capture the data modifications performed on the original data and applies them to the duplicate. This catching-up process is iterative and may need to be performed multiple times. When the shadow copy has caught up to the original copy, the data sets are switched: The copy becomes the original and the original can be deleted. In this manner, DBAs can reorganize databases with minimal impact on availability. Database reorganization is covered in more depth in [Chapter 11, “Database Performance.”](#)

But remember, most database maintenance tasks impact availability. Making backups of data, recovering data, checking data for integrity

violations, capturing database statistics, and loading new data into a database all can adversely impact availability. Tools that work in conjunction with modern storage devices to minimize or eliminate downtime also are quite useful to maintain databases while they remain online and operational. Some storage devices can make rapid snapshots of files. When database maintenance tasks can take advantage of this technique, outages can be reduced from minutes or hours to seconds.

Most database maintenance tasks impact availability.

Making changes to databases can be a major threat to availability. The impact on availability depends on the type of change and how the DBMS implements that change. When simple changes can be made using an ALTER statement, availability, though impacted, is less of a problem because this type of change can be made quickly. Changes to data objects have a greater impact on availability. The DBMS will often need to make the data unavailable while the structural definition of the database object is being changed.

As the type of change becomes more complex, the impact on availability becomes greater. DBMS vendors are adapting their products to make them more capable of nondisruptive database change, such as IBM's Database Definition on Demand initiative (see the sidebar). Certain changes require objects to be dropped and data to be deleted. Obviously, such a change causes an outage. The longer it takes to make the change, the greater the outage becomes. Using high-speed LOAD and UNLOAD utilities can shrink the duration of the outage. Automating the change, discussed next, can further diminish unavailability. Additional details of implementing database changes can be found in [Chapter 7](#).

IBM DB2: Database Definition on Demand

IBM embarked on a path of online schema evolution as of DB2 V8 with the intention of allowing DB2 databases to be altered without causing an outage. With each subsequent release of DB2 additional changes can be made without incurring an outage.

For example, some of the improved change management capabilities introduced over the past few releases of DB2 include improved ability to extend the length of columns, change the data type of columns, add

columns to indexes, and change partitioning characteristics, and improved utility support.

Over time, the Database Definition on Demand initiative will make it much easier to change DB2 databases without causing any downtime.

Automate DBA Functions

Building more automation into DBA procedures can increase overall database and application availability. When properly created, an automated task will fail less frequently than a manual task. Humans make mistakes. Computers do not make mistakes (unless a human programs them improperly). So, the more complex the task, the more it can benefit from automation.

Implementing database changes is a complex task. It stands to reason, therefore, that the automation of changes can improve availability. When an automated DBA tool that understands the DBMS and how to make changes to the database objects is used, the potential for human error is reduced. Furthermore, the time it takes for a DBA to manually generate CHANGE scripts is many times greater than the time it takes a change management tool to generate CHANGE scripts. In addition, the tool is unlikely to make errors. So, with automated database changes, less time is required to analyze the requested changes, develop the scripts to make the changes, and run the scripts to actually change the database. Overall, an automated change process is a boon to availability.

Automation of changes can improve availability.

Another task that benefits from automation is database backup and recovery. To ensure that a company can retrieve its data as quickly as possible after an outage or disaster, preplanning is necessary. Taking a proactive approach to backup and recovery management can mean the difference between a minimal outage with no data loss and a situation from which a business can never recover. Most database systems and applications provide little in the way of automated backup and recovery, nor do they provide functions that enable proactive recovery planning. DBAs need products that allow for frequent backups that exert minimal impact on the online system. An additional requirement of backup and recovery software is extremely fast recovery of data in a sequence directly related to the criticality

of the business application the data supports. Once it becomes obvious that a recovery is needed to resolve an outage, the DBA needs to ensure the proper recovery of each database object by using the backups and database logs at his disposal. This requires knowledge of the DBMS, the application, and the environment. Add-on products are available that can assist in automating recovery in a crisis situation by analyzing the situation and building the appropriate recovery scripts that will bring the system back in the shortest period of time. For an in-depth discussion of database backup and recovery, consult [Chapters 16](#) and [17](#).

Exploit High-Availability Features

If the DBMS is engineered to work with clustering and parallel technology, be sure to design databases that work well with that technology. More and more DBMS vendors are aligning their software with the capabilities of modern hardware and operating systems.

Most DBMS vendors have made additional availability features a priority due to the burgeoning need for Internet and Web support in database management systems and applications. Each new DBMS release brings additional availability options and features that can be exploited to enhance uptime and availability. Two obvious examples involve utilities and database system parameters. Running utilities and changing system parameters were tasks that traditionally required an outage. In many cases, running utilities required the database object that was being maintained to be taken offline. Similarly, the entire DBMS had to be taken down to change a system parameter. However, DBMS vendors are attacking these foes of availability. Many newer DBMS versions provide techniques to change system parameters without bringing down the DBMS (for example, the SET SYSPARM command in DB2 for z/OS). DBAs should reexamine old tasks that required database unavailability to see if those tasks can be redesigned to be performed without taking databases or the DBMS offline.

Exploit Clustering Technology

[Clustering²](#) is an option for increasing the reliability of servers. A *cluster*, quite simply, is a group of interconnected servers. The actual implementation of a server cluster can range from computers that share storage to groups of servers that can redistribute their workload from one to another with the help of special software.

Companies generally turn to clustering to improve availability and scalability. Clusters improve availability by providing alternate options in case a system fails. Clustering involves multiple, independent computing systems working together as one. So, if one of the independent systems fails, the cluster software can distribute work from the failing system to the remaining systems in the cluster. Users won't know the difference—they interact with a cluster as though it were a single server—and the resources they rely on will still be available.

Most companies consider enhanced availability to be the primary benefit of clustering. In some cases, clustering can help companies achieve “five-nines” (99.999 percent) availability. A big advantage of clustering is the ability to increase computing power by adding another server, or node, to the cluster. When throughput increases due to expanding business or publicity, systems can be kept online and available by adding more servers.

Systems can be kept online and available by adding more servers.

Clustering offers scalability benefits, too. When load exceeds the capabilities of the systems that make up the cluster, you can incrementally add more resources to the system to increase the cluster's computing power and meet processing requirements. As traffic or availability assurance increases, all or some parts of the cluster can be increased in size or number.

Another advantage is reliability. Some clusters are implemented with failover software that can reallocate the workload of one server to another when a server fails. This added reliability minimizes downtime and enhances availability. Each node in the cluster remains in contact with the others. When a node falls out of contact, the cluster recognizes the failure and initiates the failover process.

Clusters can be configured to failover in different ways. For example, when a node fails, failover can direct that processing to another node in a different location. Some configurations have an extra node in the cluster that is usually idle. When failover occurs, the idle node takes over and the cluster capacity isn't compromised. Furthermore, failover can spread the work over the other existing nodes based on capacity and throughput.

Load Balancing versus High Availability

IT professionals commonly mistake high-availability technology (such as Windows Clustering, replicated SANs, etc.) as “load balancing” technologies. *Load balancing* is a methodology to distribute workload across multiple computers (or a cluster, network links, disk drives, or other resources) to optimize resource utilization, maximize throughput, minimize response time, and avoid overload. Load balancing is usually achieved by means of dedicated software and hardware.

High availability is not about balancing. There are many ways to make databases more available, mostly determined by how much you are willing to spend. There are not many ways to make a database application load balanced on the database side of the infrastructure.

Clustering can enhance availability because failing nodes can be removed from the cluster without an outage. When the node becomes operational, it can rejoin the cluster. This may or may not require an outage, depending on the clustering implementation.

Given all these advantages, why wouldn't every IT organization choose a clustering configuration? Of course, the primary reason is cost. A cluster requires multiple machines—and one machine is always cheaper than multiple machines, at least initially. Another consideration is that applications may need to be modified to enable failover, depending on the cluster implementation.

Most operating systems support some degree of clustering. Windows servers can failover if one goes down. A cluster of IBM mainframes can completely share system resources and appear for all intents and purposes as one system. Furthermore, from a database perspective, the DBMS software needs to be programmed to understand and take advantage of the clustering support in the operating systems. Some DBMS products do this better than others.

Additionally, clustering is useful for masking the impact of routine maintenance. When a server node needs to be taken offline for maintenance, its work can be shifted to another server. For example, if memory needs to be added to the motherboard of a server, that server must be shut down. If it participates in a cluster, its workload can be routed to other nodes in the cluster while the maintenance is being performed, resulting in an outage-less maintenance task.

Clustering is useful for masking the impact of routine maintenance.

Types of Clustering

Shared-disk and shared-nothing architectures are the predominant approaches to clustering. The names are fairly accurate descriptions of each type. [Chapter 2](#) contains an overview of these types of clustering.

Essentially, shared-nothing and shared-disk offer two differing techniques for clustering. Shared-nothing is the predominant clustering architecture used by most computing and database implementations. Scalability and performance are hallmarks of shared-nothing clustering and make it ideal for analytical and data warehousing applications. Shared-disk clustering requires a controlling facility to manage access to the shared data. Shared-disk clustering is ideal for handling data warehousing workloads because their read-only nature does not require locking and therefore will not stress a shared-disk system in the same way that OLTP workloads can.

A Few Database Examples

An Oracle standby database is a simple example of one type of cluster. The primary database is mirrored to the standby database, which can step in and take over in the event of a failure. The redundancy can be expensive but is cost-effective for businesses that cannot afford long outages.

Another clustering-related, high-availability feature of Oracle is its Real Application Clusters (RAC) support. RAC provides for the transparent deployment of a single database across pools of servers, providing fault tolerance from hardware failures or planned outages. Oracle RAC provides a single image installation and management. Database administrators have a single point of control to install and manage an Oracle RAC cluster using provided graphical user interface (GUI) or command-line tools.

Oracle RAC eliminates an individual server as a single point of failure. If a node in a server pool fails, the Oracle Database continues to run on the remaining servers in the pool. Individual servers can be shut down for maintenance while application users continue to work.

Another example of a clustered system is IBM's Sysplex multiprocessor line, which splits tasks among parallel processors. Individually, the processors are less powerful than traditional bipolar predecessors, but

combined, they crunch data faster by assigning work to open processors rather than requiring users to wait for cycles on a single processor. DB2 for z/OS can be set up to take advantage of this type of parallel processing. The sidebar “[Data Sharing and DB2 for z/OS](#)” offers more details.

Keep in mind that standard database maintenance software may not run very efficiently on clustered and parallel systems. To reduce costs and improve availability, the tools used by DBAs must be capable of understanding and exploiting the clustering and parallel technologies being used. Otherwise, products that run slowly and inefficiently because they were built for a different hardware environment will negate the benefits of the parallel technology. For example, to reorganize databases and handle backup and recovery functions for parallel processing environments, DBAs need maintenance utilities written specifically to take advantage of parallel processors.

Data Sharing and DB2 for z/OS

DB2 Data Sharing allows applications running on multiple DB2 subsystems to concurrently read and write to the same data sets. Simply stated, Data Sharing enables multiple DB2 subsystems to behave as one.

The primary benefit of Data Sharing is to provide increased availability to data. With Data Sharing, data is available for direct access across multiple DB2 subsystems. Furthermore, applications can be run on multiple smaller, more competitively priced microprocessor-based machines, thereby enhancing data availability and the price/performance ratio.

An additional benefit is expanded capacity. Capacity is increased because more processors are available to execute the DB2 application programs. Instead of a single DB2 subsystem on a single logical partition, multiple central processor complexes (CPCs) can be used to execute a program (or even a single query).

DB2 Data Sharing requires an IBM Parallel Sysplex. An IBM Sysplex is a set of z/OS (mainframe) images³ that are connected and coupled by sharing one or more Sysplex timers. A Parallel Sysplex is a basic Sysplex that additionally shares a coupling facility. The coupling facility provides external shared memory and a set of hardware protocols that allow enabled applications and subsystems to share data with integrity

by using external shared memory. A Parallel Sysplex enhances scalability by extending the ability to increase the number of processors within a single z/OS image with the ability to have multiple z/OS images capable of cooperatively processing a shared workload.

For DB2 on non-mainframe platforms, the DB2 pureScale architecture is similar to DB2 for z/OS Data Sharing. IBM created the pureScale architecture based on mainframe Data Sharing concepts and integrated that together with the most current distributed technologies to deliver availability and scalability services to distributed platforms.

Database Architecture and NoSQL

For certain data-intensive applications—such as indexing a large number of documents, serving pages on high-traffic Web sites, and delivering streaming media—traditional relational database systems may not be the best choice for delivering high availability and performance. NoSQL database systems are designed to work well in such circumstances and are often associated with “Big Data” applications.

Although the NoSQL movement is gaining popularity, there isn’t exactly much rigor in terms of defining exactly what a NoSQL database system is, or what it must be able to do. At a high level, NoSQL implies nonrelational, distributed, flexible, and scalable. Many NoSQL database products are also open source. Common attributes of a NoSQL database system include the lack of a schema, simplicity of use, replication support, and an “eventually consistent” capability (instead of the typical ACID transaction capability).

NoSQL databases have been deployed in the real world at Digg (a 3TB implementation for green badges) and at Facebook (a 50TB implementation for in-box search). Twitter also uses NoSQL, having moved from its initial MySQL foundation to the Java-based NoSQL Cassandra DBMS.

Summary

Organizations have to find a balance between the seemingly incompatible needs for 24/7 uptime and periodic maintenance. A poorly maintained database is a business inhibitor and will be nearly impossible to restore in the event of a crisis. There are alternatives to the native database utilities that can deliver maintenance and backup functionality while providing continuous

availability of the database and associated applications. In many instances, critical applications directly affect revenue. Thus, the DBA must implement a maintenance and backup strategy that provides optimum availability.

Additionally, the DBA must remain alert to all of the potential problems that can cause a lack of availability. These causes run the gamut from hardware problems to software bugs to human error. Obviously, each type of database unavailability has different effects on the organization, the users, and the DBA. Although each involves an outage from the end user perspective, some are easier for the DBA to cope with than others. At any rate, the DBA must be prepared to resolve any type of availability problem that impacts a database user's ability to access and modify data. This is a complex and challenging task because its scope is truly enterprise-wide—just about every component of the IT infrastructure can impact database availability.

Remain alert to all of the potential problems that can cause a lack of availability.

However, if the DBA has the right skills and training, prudent DBA practices and 24/7 availability need not be mutually exclusive. It just takes the right tools, a little planning, and a lot of diligence.

Review

1. Define what is meant by the term [availability](#).
2. Compare and contrast performance problems with availability problems.
3. Describe three types of technology that can help to alleviate availability problems.
4. Hardware and software failures are the only causes of data unavailability: true or false?
5. What impact does corrupt and invalid data have on availability?
6. What factors contribute to the cost of downtime?
7. What is the typical cause of an unplanned outage?
8. What percentage of downtime is caused by planned outages?
9. What are nondisruptive database utilities, and why are they important

for maintaining database availability?

10. What are the four (high-level) steps that can be taken to improve database availability?

Suggested Reading

Bertucci, Paul. *Microsoft SQL Server High Availability*. Indianapolis, IN: SAMS (2005). ISBN 0-672-32625-6.

Devraj, Venkat S. *Oracle 24x7: Real-World Approaches to Ensuring Database Availability*. Berkeley, CA: Osborne/McGraw-Hill (2000). ISBN 0-07-211999-3.

Eaton, Chris. *High Availability Guide for DB2*. Upper Saddle River, NJ: Prentice Hall (2004). ISBN 978-0-7686-8220-5.

Jesse, Scott, et al. *Oracle Database 11g Release 2 High Availability: Maximize Your Availability with Grid Infrastructure, RAC and Data Guard*. Berkeley, CA: McGraw-Hill (2011). ISBN 978-0-07-175208-4.

O'Reilly Media. *Big Data Now*. Sebastopol, CA: O'Reilly (2011). ISBN 978-1-449-31518-4.

Otey, Michael. *Microsoft SQL Server 2008 High Availability with Clustering & Database Mirroring*. Berkeley, CA: McGraw-Hill (2010). ISBN 978-0-07-149813-5.

Piedad, Floyd, and Michael Hawkins. *High Availability: Design, Techniques and Processes*. Upper Saddle River, NJ: Prentice Hall (2001). ISBN 0-13-096288-0.

9. Performance Management

When non-DBAs think about what it is that a DBA does, performance monitoring and tuning are quite frequently the first tasks that come to mind. This should not be surprising. Almost anyone who has come in contact with a computer has experienced some type of performance problem. Moreover, relational database systems have a notorious reputation (mostly undeserved) for poor performance.

This chapter, as well as the following three, will discuss performance monitoring, tuning, and management within the context of database administration. [Chapter 9](#) defines performance, discusses the difference between performance monitoring and performance management, looks at managing service levels, and defines three specific subsets of database performance management. [Chapters 10](#) through [12](#) delve further into the three subsets.

Defining Performance

Most organizations monitor and tune the performance of their IT infrastructure. This infrastructure encompasses servers, networks, applications, desktops, and databases. However, the performance management steps taken are usually reactive. A user calls with a response-time problem. A tablespace runs out of disk storage space in which to expand. The batch window extends into the day. Someone submitted a “query from hell” that just won’t stop running. Those of you in the trenches can relate—you’ve been there, done that.

Performance management is usually reactive.

Handling performance problems is truly an enterprise-wide endeavor. However, the task of enterprise performance management frequently becomes the job of the DBA group. Anyone who has worked as a DBA for any length of time knows that the DBMS is usually “guilty until proven innocent.” Every performance problem gets blamed on the database regardless of its true cause. DBAs need to be able to research and ascertain the source of all performance degradation, if only to prove that it is not caused by a database problem. Thus, DBAs must be able to understand at

least the basics of the entire IT infrastructure, but they also need to have many friends who are experts in other related fields (such as networking, operating systems, and communication protocols). Possessing a sound understanding of the IT infrastructure enables DBAs to respond effectively when performance problems arise. Event-driven tools exist on the market that can make performance management easier by automatically invoking predefined actions when specific alerts are triggered. For example, an alert can be set to proactively reorganize a database when it reaches its storage capacity or to allocate more memory when the DBMS is reaching its limit. Moreover, other tools exist that can ease the burden of performance management and analysis. However, many of the supposedly proactive steps taken against completed applications in production are truly mostly reactive. Let's face it, DBAs are often too busy taking care of the day-to-day tactical database administration tasks to proactively monitor and tune their systems to the degree they wish they could.

Handling performance problems is truly an enterprise-wide endeavor.

All of this discussion is useful, but it begs the question: Just what do we mean by *database performance*? You need a firm definition of database performance before you can plan for efficiency. Think, for a moment, of database performance using the familiar concepts of supply and demand. Users request information from the database. The DBMS supplies information to those requesting it. The rate at which the DBMS supplies the demand for information can be termed *database performance*. However, this definition captures database performance only in a most simplistic form.

We need a more comprehensive definition of database performance. Five factors influence database performance: workload, throughput, resources, optimization, and contention.

The *workload* is a combination of online transactions, batch jobs, ad hoc queries, data warehousing analysis, and system commands directed through the system at any given time. Workload can fluctuate drastically from day to day, hour to hour, and even minute to minute. Sometimes workload is predictable (such as heavy month-end processing of payroll, or very light access after 7:00 p.m., when most users have left for the day), whereas workload is very unpredictable at other times. The overall workload has a

major impact on database performance.

Throughput defines the overall capability of the computer to process data. It is a composite of I/O speed, CPU speed, parallel capabilities of the machine, and the efficiency of the operating system and system software. The hardware and software tools at the disposal of the system are known as the *resources* of the system. Examples of resources include the database kernel, disk storage devices, random access memory chips, cache controllers, and microcode.

The fourth defining element of database performance is *optimization*. All types of systems can be optimized, but relational databases are unique in that query optimization is primarily accomplished internal to the DBMS. However, many other factors need to be optimized (such as SQL formulation and database parameters) to enable the database optimizer to create the most efficient access paths.

When the demand (workload) for a particular resource is high, contention can result. [*Contention*](#) is the condition where two or more components of the workload are attempting to use a single resource in a conflicting way (e.g., dual updates to the same piece of data). As contention increases, throughput decreases.

Therefore, database performance can be defined as the optimization of resource use to increase throughput and minimize contention, enabling the largest possible workload to be processed. Of course, I do not advocate managing database performance in a vacuum. Applications regularly communicate with other subsystems and components of the IT infrastructure. Each of these must also be factored into the overall performance planning of your organization. However, it is wise to place limits on the DBA's actual responsibility for performance tuning outside the scope of this definition. If the task is not included in the definition above, it probably requires expertise outside the scope of database administration. Therefore, performance management tasks not covered by this description should be handled by someone other than the DBA—or at least shared among the DBA and other technicians.

Database performance is the optimization of resource use to increase throughput and minimize contention, enabling the largest possible workload to be processed.

A Basic Database Performance Road Map

Planning for database performance management is a crucial component of any application implementation. Therefore, the DBA needs to forge a basic plan to ensure that database performance management and analysis is accomplished for all database applications across the organization. A complete performance management plan will include tools to help monitor application performance and tune the database and SQL.

Following the 80/20 rule (see the sidebar), the first step should be to identify the most troublesome areas. However, this is not always as easy as it might seem.

The most likely culprit for most database application performance problems is inefficient SQL and application code. In my experience, 75 to 80 percent of all database performance problems can be traced to poorly coded SQL or application logic. This does not mean that the SQL in applications is necessarily bad to begin with. Although an application may be 100 percent tuned for rapid relational access when it first moves into the production environment, it can suffer performance degradation over time. This degradation can occur for many reasons, such as database growth, new data access patterns, additional users, changes in the business, and so on.

The most likely culprit for most performance problems is inefficient SQL and application code.

The 80/20 Rule

The *80/20 rule*, also known as the *Pareto Principle*, is an old maxim stating that 80 percent of the results come from 20 percent of the effort. This rule is usually applicable to most efforts. Whether or not the percentages are precisely 80 percent and 20 percent, the underlying logic of the rule holds—namely, that a small amount of effort brings the most rewards.

For example, the Pareto Principle as applied to database performance tuning can be stated as 80 percent of the results of tuning come from 20 percent of the effort. Additionally, it can just as easily be applied to mean that 20 percent of your database applications cause 80 percent of your problems.

The bottom line is that, from the perspective of database performance tuning, the wise DBA will concentrate on the most likely causes of performance problems first, because he will receive a high return on his tuning investment.

Of course, the SQL and application code can be just plain bad to begin with, too. Any number of problems can cause poorly performing SQL, including

- Table scans
- Lack of appropriate indexes
- Improper indexing choices
- Not using the available indexes
- Outdated database statistics
- Tables joined in a suboptimal order
- Application joins instead of (usually) more efficient SQL joins
- Improper join method (nested loop, merge scan, etc.)
- Efficient SQL inside of inefficient application code (loops)
- Inefficient subquery formulation (exists, not exists, etc.)
- Unnecessary sorting (distinct, group by, order by, union)

Finding the SQL statements that are the most expensive in a large shop is an extremely difficult thing to do.¹ Resource-hogging SQL statements might be hiding in one of hundreds or even thousands of programs. Interactive users who produce dynamic, ad hoc SQL statements might reside anywhere, and any one person who is generating ad hoc queries can severely affect overall production performance.

A good approach is to use an SQL monitor that identifies all SQL running anywhere in your environment. Typically, these tools rank SQL statements based on the amount of resources being consumed and track a statement back to who issued it and from what program. Once you have identified the top resource-consuming statements, you can concentrate your tuning efforts on the most costly statements.

However, it is not always obvious how to tune poorly coded SQL statements. The proper coding and tuning of SQL statements is a detailed

endeavor. In-depth strategies for SQL tuning and additional application performance management details can be found in [Chapter 12](#), “[Application Performance](#).”

The proper coding and tuning of SQL statements is a detailed endeavor.

Of course, other factors can negatively impact database performance. It is wise to periodically check the overall performance of the database instance and the server operating system. Some quick items to check include the following:

- Memory allocation (buffer/cache for data, SQL, authorization)
- Logging options (log cache, log size, Oracle rollback segments)
- I/O efficiency (separation of tables and indexes on disk, database size, fragmented and extended files)
- Overall application and database workload on the server
- Database schema definitions

To assure optimum database performance, plan on combining a good definition of database performance with a detailed performance plan specific to your shop.

Monitoring versus Management

Unfortunately, the DBA usually attacks performance in a reactive manner. A user calls with a response-time problem. A database runs out of space. The batch window extends into the day. The problem has happened and now it needs to be remedied. Such activity is purely reactive.

Even many of the supposedly proactive steps taken against completed production applications might be considered reactive. A change to a completed application that requires code to be rewritten cannot reasonably be considered proactive. A proactive approach would have involved correcting the problem before completing the application.

Some event-driven tools can be used to make performance tuning easier by automatically taking predefined actions when prespecified alerts are triggered. This is the first step toward performance management. Managing performance differs from monitoring performance because it combines

monitoring with a detailed plan for resolving problems when they arise.

Performance management consists of three specific components that need to be performed in conjunction with each other: monitoring, analysis, and correction, as shown in [Figure 9.1](#). *Monitoring* is the first component of performance management. It consists of scanning the environment, reviewing the output of instrumentation facilities, and generally watching the system as it runs. Monitoring is the process of identifying problems.

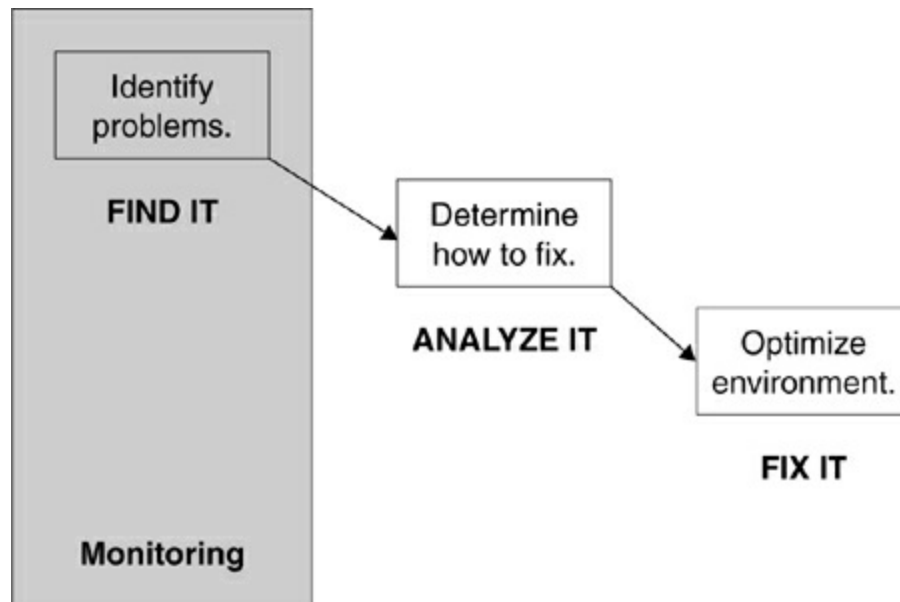


Figure 9.1. The components of performance management

Performance management consists of three steps: monitoring, analysis, and correction.

Analysis is the second component of performance management. A monitoring task can generate hundreds or thousands of messages, or reams and reams of paper reports. A monitor collects the pertinent information for making performance tuning and optimization decisions, but it is essentially dumb. A monitor cannot independently make decisions based on the information it has collected. This requires analysis—and analysis typically is performed by a skilled technician such as a DBA.

Optimization—the corrective action—is the third component of performance management. Some performance tools allow the technician to automate certain aspects of performance management by automatically kicking off corrective actions when the monitoring agent identifies certain

conditions. However, most of these tools are limited in scope. Furthermore, a skilled technician must set up the automation to make sure that the appropriate optimization step is taken at the correct time. Eventually performance management tools and solutions will become intelligent—with built-in knowledge of how to optimize a DBMS and with the ability to learn what works best from tuning exercises.

Performance management can be achieved only by using a proactive performance plan. Many problems can be identified and solutions mapped out in advance, before an emergency occurs. With a proper plan, correction of performance problems becomes easier, and indeed, some performance problems can be avoided altogether because a potential problem-causing situation can be corrected before a problem occurs.

Performance management can be achieved only by using a proactive performance plan.

For true proactive performance management to be achieved, the DBA must plan the performance of an application before it is completed. This requires the DBA to be involved in the application development life cycle and to ensure that performance is designed into the application. Achieving such a high degree of DBA involvement in the application development process can be problematic. DBAs are pressed for time, and it seems as if there is always time to do it over later instead of simply doing it right the first time.

Reactive versus Proactive

Reactive performance management will always be required because unplanned performance problems will always occur. It is impossible to foresee every type of performance problem; after all, systems and applications change over time. Reactive performance management is not, in itself, a bad thing, but it is a manual, time-consuming process. *Proactive performance management* combines forethought, planning, and automation to minimize reactive monitoring and tuning. In other words, proactive performance management reduces the amount of time, effort, and human error involved in implementing and maintaining efficient database systems.

Proactive performance management combines forethought, planning, and automation to minimize reactive monitoring and tuning.

Although proactive performance management has a greater payback than reactive, many DBA organizations struggle to implement a plan for proactively attacking performance problems. In most cases this is because the DBAs are overburdened with reactive performance chores. Indeed, when the boss is breathing down their neck, the phone keeps ringing with new problems, and a queue of developers are lined up outside the DBA's office, it is unreasonable to expect the DBA to proactively seek out additional work.

But keep in mind that many of the problems that require reactive attention could be avoided with a proper proactive performance management approach.

Preproduction Performance Estimation

Ideally, DBAs and developers should engineer high performance into their applications during design and construction by implementing a methodology. Such a methodology must address the ADLC by incorporating tactics to achieve high performance into the creation of the database and application code. A rigorous process focusing on verifiable results can build performance into applications and databases, thereby eliminating costly redesign and recoding efforts—at least with respect to most performance problems.

Problems identified earlier in the ADLC are easier to fix and cost less to fix than problems identified later in the application's life, as shown in [Figure 9.2](#). Proactive performance management can reduce the cost of application development because it occurs before the application becomes operational in a production environment. Correcting problems after an application is operational is the costliest method of performance management, because users are relying on the operational application to perform their jobs. Performance problems in a production application can increase the amount of time it takes to perform mission-critical work, such as servicing customers. Moreover, severe performance problems can cause outages, as discussed in the preceding chapter.

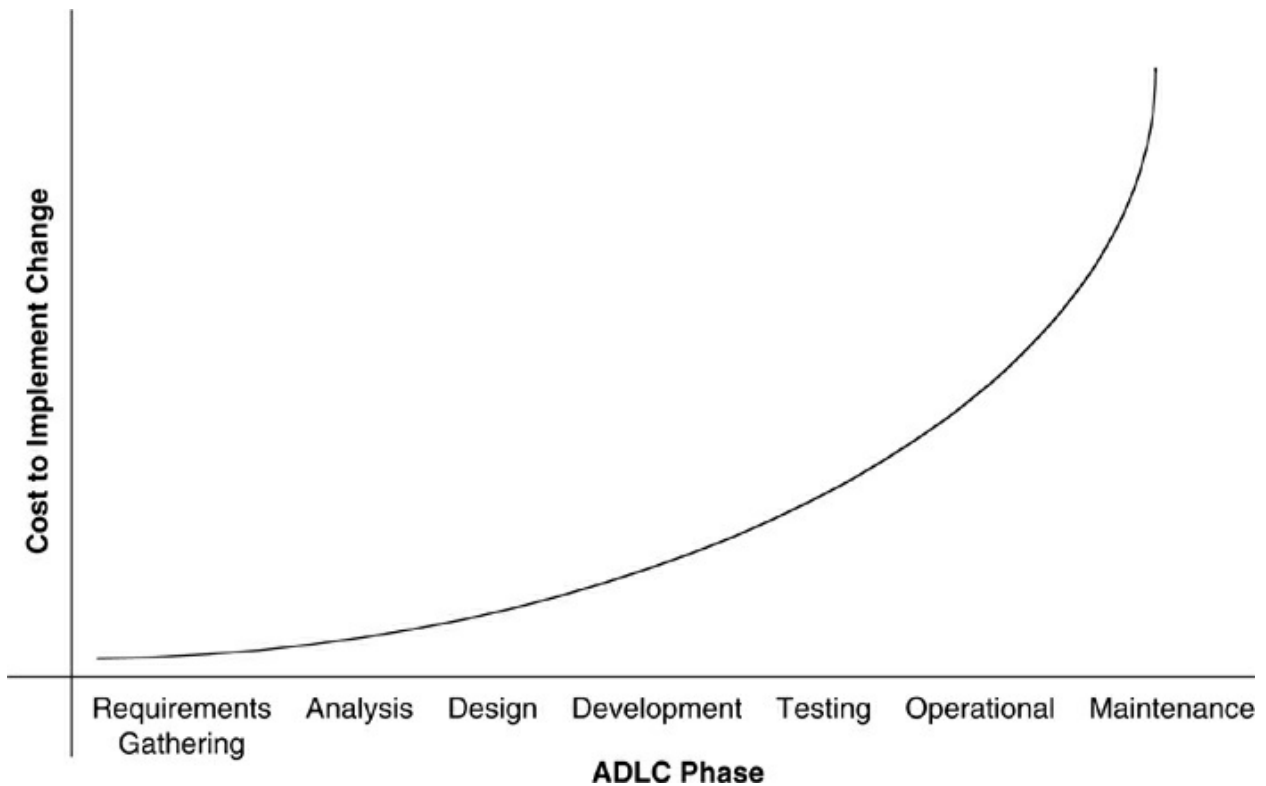


Figure 9.2. Cost of performance problems across the application development life cycle

Estimating the performance of applications differs from analyzing and optimizing single database queries. Performance should be modeled for the entire application, because individual queries may optimize at the expense of other queries. A model will show the overall effect of all the queries and how they affect each other's performance. Such a model enables the DBA to optimize overall performance.

Performance should be modeled for the entire application.

Creating an accurate performance model is an iterative process. Each change must be reviewed and updated and its impact gauged for effectiveness. DBAs, SAs, application developers, and capacity planners must cooperate to share information and address any business requirement issues that may affect the performance criteria.

Historical Trending

Capturing and analyzing resource usage trends and performance statistics over time is another valuable performance task. Historical performance and

resource trending allows DBAs to predict the need for hardware upgrades weeks, and perhaps months, in advance. Administrators can track key performance statistics (such as buffer hit ratios, file I/O, and log switches) and store that information in *tracker tables* in the database. This provides valuable historical information that can be reported on and analyzed. DBAs can track performance and resource consumption and predict when hardware resources will be consumed by increasing usage. Furthermore, historical trends can illuminate periods when database performance is slower than usual due to increased user activity.² For example, database applications tend to run slower the first three days of the month due to month-end processing requirements. Maintaining key historical performance indicators can provide a huge benefit to DBAs as they attempt to comprehend the performance characteristics of their applications, databases, and systems.

Maintaining key historical performance indicators can provide a huge benefit to DBAs.

Service-Level Management

Service-level management (SLM) is the “disciplined, proactive methodology and procedures used to ensure that adequate levels of service are delivered to all IT users in accordance with business priorities and at acceptable cost.”³ In order to effectively manage service levels, a business must prioritize its applications and identify the amount of time, effort, and capital that can be expended to deliver service for those applications.

A service level is a measure of operational behavior. SLM ensures that applications behave accordingly by applying resources to those applications based on their importance to the organization. Depending on the needs of the organization, SLM can focus on availability, performance, or both. In terms of availability, the service level might be defined as “99.95 percent uptime from 9:00 a.m. to 10:00 p.m. on weekdays.” Of course, a service level can be more specific, stating that “average response time for transactions will be 2 seconds or less for workloads of 500 or fewer users.”

SLM ensures that resources are applied to applications based on their importance to the organization.

For an SLA to be successful, all parties involved must agree on stated

objectives for availability and performance. The end users must be satisfied with the performance of their applications, and the DBAs and technicians must be content with their ability to manage the system to the objectives. Compromise is essential to reach a useful SLA.

In practice, though, many organizations do not institutionalize SLM. When new applications are delivered, there may be vague requirements and promises of subsecond response time, but the prioritization and budgeting required to assure such service levels are rarely tackled unless the IT function is outsourced. Internal IT organizations are loath to sign SLAs because any SLA worth pursuing will be difficult to achieve. Furthermore, once the difficulties of negotiating an SLA are completed, the business could very well turn around and outsource the SLA to a lower-cost provider than the internal IT group.

But do not misunderstand. The failure of SLM within most businesses lies with both IT organizations and business users. The business users frequently desire better service but are not willing to make the effort to prioritize their needs correctly or to pay additional cash to achieve better service.

Another potential problem with SLM is the context of the service being discussed. Most IT professionals view service levels on an element-by-element basis. In other words, the DBA views performance based on the DBMS, the SA views performance based on the operating system or the transaction processing system, and so on. SLM properly views service for an entire application. However, it can be difficult to assign responsibility within the typical IT structure. IT usually operates as a group of silos that do not work together very well. Frequently, the application teams operate independently from the DBAs, who operate independently from the SAs, as shown in [Figure 9.3](#). When an application team has staffed an application DBA function, that team may not communicate effectively with the corporate DBA silo. These fractured silos make cooperation toward a common application service level difficult.

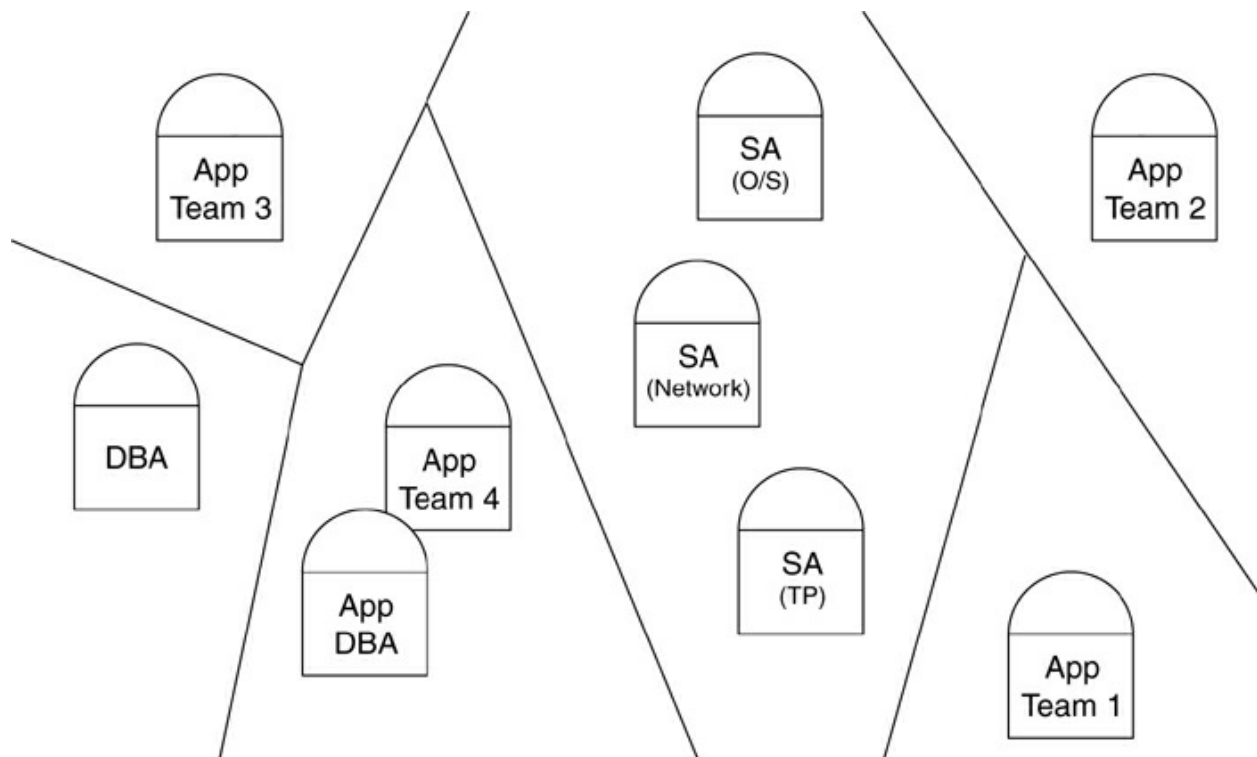


Figure 9.3. IT silos in a fractured environment

To achieve end-to-end SLM, these silos need to be broken down. The various departments within the IT infrastructure need to communicate effectively and cooperate with one another. Failing this, end-to-end SLM will be difficult, if not impossible, to implement.

SLM is a beneficial practice: A robust SLM discipline makes performance management predictable. SLM manages the expectations of all involved. Without an SLA, how will the DBA and the end users know whether an application is performing adequately? Not every application can, or needs to, deliver subsecond response time. Without an SLA, business users and DBAs may have different expectations, resulting in unsatisfied business executives and frustrated DBAs—not a good situation.

A robust SLM discipline makes performance management predictable.

With SLM in place, DBAs can adjust resources by applying them to the most-mission-critical applications as defined in the SLA. Costs will be controlled and capital will be expended on the portions of the business that are most important to the business. Without SLM in place, an acceptable

performance environment will be ever elusive. Think about it; without an SLA in place, if the end user calls up and complains to the DBA about poor performance, there is no way to measure the veracity of the claim or to gauge the possibility of improvement within the allotted budget.

Types of Performance Tuning

A database application requires constant interaction between disparate computing resources in order to operate efficiently and according to specifications. Realistically, though, the tuning of a database application can be broken down into three components: system tuning, database tuning, and application tuning. Indeed, all these areas are related, and certain aspects of tuning require an integrated approach. However, for clarity, we'll discuss these areas separately.

System Tuning

System tuning occurs at the highest level and has the greatest impact on the overall health of database applications because every application depends on the system. For the purposes of this discussion, we will define the system as comprising the DBMS itself and all of the related components on which it relies. No amount of tuning is going to help a database or application when the server it is running on is short on resources or improperly installed.

The DBMS can and must be tuned to assure optimum performance. The way in which the DBMS software is installed, its memory, disk, CPU, other resources, and any configuration options can impact database application performance.

The DBMS can and must be tuned to assure optimum performance.

The other systems software with which the DBMS interacts includes the operating system, networking software, message queuing systems, middleware, and transaction processors. System tuning comprises installation, configuration, and integration issues, as well as ensuring connectivity of the software to the DBMS and database applications.

Database Tuning

Performance can be impacted by the physical design of the database,

including normalization, disk storage, number of tables, index design, and use of DDL and its associated parameters. The physical location of database files on disk systems will have an impact on the performance of applications accessing the data. As more data is stored on the same disk device, the possibility of performance degradation increases.

The physical location of database files on disk systems impacts the performance of applications accessing the data.

However, design is not the only component of database performance. The organization of the database will change over time. As data is inserted, updated, and deleted from the database, the efficiency of the database will degrade. Moreover, the files that hold the data may need to expand as more data is added. Perhaps additional files, or file extents, will need to be allocated. Both disorganization and file growth can degrade performance.

Indexes also need to be monitored, analyzed, and tuned to optimize data access and to ensure that they are not having a negative impact on data modification.

Application Tuning

The application itself must be designed appropriately and monitored for efficiency. Most experts agree that as much as 75 percent of performance problems are caused by improperly coded applications. SQL is the primary culprit; coding efficient SQL statements can be complicated. Developers need to be taught how to properly formulate, monitor, and tune SQL statements.

However, not all application problems are due to improperly coded SQL. The host language application code in which the SQL has been embedded may be causing the problem. For example, Java, COBOL, C++, Ruby, or Visual Basic code may be inefficient, causing database application performance to suffer.

Not all application problems are due to improperly coded SQL.

Managing the performance for each of these three areas—the system, the database, and the application—will be discussed in depth in [Chapters 10, 11, and 12](#).

Performance Tuning Tools

Database tools are helpful to effectively manage database performance. Some DBMS vendors provide embedded options and bundled tools to address database performance management. However, these tools are frequently insufficient for large-scale or heavily used database applications. Fortunately, many third-party tools will effectively manage the performance of mission-critical database applications. Tools that enable DBAs to tune databases fall into two major categories: performance management and performance optimization.

Many third-party tools can effectively manage the performance of mission-critical database applications.

Many different types of performance management tools are available:

- *Performance monitors* enable DBAs and performance analysts to gauge the performance of applications that access databases in one of three ways: real time, near real time (intervals), or based on historical trends. The more advanced performance monitors are agent based.
- *Performance estimation tools* provide predictive performance estimation for entire programs and SQL statements based on access paths, operating environment, and a rules or inference engine.
- *Capacity planning tools* enable DBAs to analyze the current environment and database design and perform “what-if” scenarios on both.
- *SQL analysis and tuning tools* provide graphical and/or textual descriptions of query access paths as determined by the relational optimizer. These tools can execute against single SQL statements or entire programs.
- *Advisory tools* augment SQL analysis and tuning tools by providing a knowledge base of tips on how to reformulate SQL for optimal performance. Advanced tools may automatically change the SQL (on request) based on the coding tips in the knowledge base.
- *System analysis and tuning tools* enable the DBA to view and change database and system parameters using a graphical interface (e.g., cache and/or buffer pool tuning, log sizing).

In the performance optimization category, several tools can be used to tune databases:

- *Reorganization tools* automate the process of rebuilding optimally organized databases. Databases can cause performance problems due to their internal organization (e.g., fragmentation, row ordering, storage allocation).
- *Caching tools* work to buffer frequently used data in memory, which can be accessed faster than secondary disk storage. These tools can augment the performance of the DBMS cache or, more commonly, integrate with the disk storage subsystem.
- *Compression tools* enable DBAs to minimize the amount of disk storage used by databases, thereby reducing overall disk utilization and, possibly, elapsed query/program execution time, because fewer I/Os may be required. (Caution: Compression tools can also increase CPU consumption due to the overhead of their compress/decompress algorithms.)
- *Sorting tools* can be used to sort data prior to loading databases to ensure that rows will be in a predetermined sequence. Additionally, sorting tools can be used in place of ORDER BY or GROUP BY SQL. Retrieving rows from a relational database is sometimes more efficient using SQL and ORDER BY rather than SQL alone followed by a stand-alone sort of the SQL results set.

The DBA will often need to use these tools in conjunction with one another—integrated and accessible from a central management console. This enables the DBA to perform core performance-oriented and database administration tasks from a single platform.

The DBA will often need to use these tools in conjunction with one another.

Many DBMS vendors provide solutions to manage their databases only; for example, Oracle provides Oracle Enterprise Manager, IBM offers Data Studio for DB2, and Microsoft provides SQL Server Management Studio for this purpose. Third-party vendors provide more robust options that act across heterogeneous environments such as multiple different database servers or operating systems. One example is Embarcadero Technologies' DBArtisan.

In general, it is wise to use the DBMS vendor solution as your only management tool if your shop has just a single DBMS. Organizations with multiple DBMS engines running across multiple operating systems should investigate the third-party tool vendors with heterogeneous support (perhaps in addition to the single-solution tools).

[Chapter 23](#), “[DBA Tools](#),” provides additional information on tools that simplify database administration and performance management tasks.

DBMS Performance Basics

We have defined database performance and discussed it from a high level. Before we delve into the specifics of system, database, and application performance, let’s examine some rules of thumb for achieving your DBMS-related performance goals:

- *Do not over-tune.* Most DBAs are more than happy to roll up their sleeves and get their hands dirty with the minute technical details of the DBMS. Sometimes this is required. However, as a DBA, you should always keep in mind the business objectives of the databases and applications you manage. It is wise to manage performance based on the expectations and budget of the business users. Even though it might be an interesting intellectual challenge for you to fine-tune a query to its best performance, doing so may take too much time away from your other duties. It is best to stop tuning when performance reaches a predefined service level for which the business users are willing to pay.
- *Remain focused.* As a DBA, you should understand the goal for each task you perform and remain focused on it. This is important because the DBMS is complex, and when you are tuning one area you might find problems in another. If so, it is best to document what you found for later and continue with the tuning task at hand. Furthermore, by jumping around trying to tune multiple things at once, you will have no idea of each task’s impact on the environment.
- *Do not panic.* DBAs are expected to know everything about the DBMS they manage. However, this is an unreasonable expectation. “I don’t know, but I’ll find out” is one of the most important sentences in your communications arsenal. A good DBA knows where to look for answers and whom to call for help.
- *Communicate clearly.* Communication is key to assuring properly

tuned, high-performance database systems. The DBA must be at the center of that communication, coordinating discussions and workload among the business users, programmers, managers, and SAs. Furthermore, the world of IT in general, and database technology in particular, sometimes uses a language all its own. Many similar and confusing terms are thrown about, and folks are expected to understand what they mean. Be sure to clearly define even basic terms so that you're all speaking the same language.

- *Accept reality.* Many organizations talk about being proactive but in reality have very little interest in stopping performance problems before they happen. Yet, every organization is interested in fixing performance problems after the problems occur. This can be a frustrating environment for the DBA, who would rather set up preventive maintenance for the DBMS environment. Alas, this requires budget, time, and effort—all of which are in short supply for strapped IT organizations. As a DBA, you must sometimes be content to accept reality and deal with problems as they occur—even when you know there are better ways of tackling performance management.

Summary

Applications that access relational databases are only as good as the performance they achieve. The wise organization will implement a comprehensive performance monitoring, tuning, and management environment consisting of policies, procedures, and integrated performance management tools and utilities.

Applications that access relational databases are only as good as the performance they achieve.

Review

1. Define *database performance*.
2. During which phase of the application development life cycle is it most costly to make a change?
3. What are the three factors that need to be addressed in database application tuning?
4. What does the 80/20 rule mean and how should it be applied to

database application performance tuning?

5. Name and describe the three steps of database performance management.
6. Discuss the problems involved in implementing service-level management for database applications.
7. What two categories of database tools can be used to help deliver database performance management?
8. What percentage of performance problems do most experts attribute to improperly coded database applications?
9. What is the condition in which two or more components of the workload are attempting to use a single resource in a conflicting way?
10. Compare and contrast reactive and proactive database performance management.

Bonus Question

Why is it a good idea to capture and store historical performance statistics?

What types of questions can historical statistics help a DBA to answer?

Suggested Reading

Dunham, Jeff. *Database Performance Tuning Handbook*. New York, NY: McGraw-Hill (1998). ISBN 0-07-018244-2.

Gunther, Neil J. *The Practical Performance Analyst*. Lincoln, NE: Authors Choice Press (2000). ISBN 0-595-12674-X.

Harrison, Guy. *Oracle Performance Survival Guide: A Systematic Approach to Database Optimization*. Upper Saddle River, NJ: Prentice Hall (2009). ISBN 978-0-13-701195-7.

Koch, Richard. *The 80/20 Principle*. New York, NY: Currency/Doubleday (1998). ISBN 0-385-49170-0.

Lawson, Christopher. *The Art and Science of Oracle Performance Tuning*. Birmingham, UK: Curlingstone (2003). ISBN 1-904347-01-0.

Loosley, Chris, and Frank Douglas. *High-Performance Client/Server*. New York, NY: John Wiley & Sons (1998). ISBN 0-471-16269-8.

Mitra, Sitansu S. *Database Performance Tuning and Optimization*. New York, NY: Springer-Verlag (2003). ISBN 0-387-95393-0.

Mullins, Craig S. "Dissecting Database Performance." In *Database Trends & Applications* (October 2001). http://craigsmullins.com/dbta_002.htm.

Niemic, Richard. *Oracle Database 11g Release 2 Performance Tuning Tips & Techniques*. New York, NY: Oracle Press (2012). ISBN 978-0-07-178027-8.

Shasha, Dennis A. *Database Tuning: A Principled Approach*. Englewood Cliffs, NJ: Prentice Hall (1992). ISBN 0-13-205246-6.

Sturm, Rick, Wayne Morris, and Mary Jander. *Foundations of Service Level Management*. Indianapolis, IN: SAMS Publishing (2000). ISBN 0-672-31743-5.

10. System Performance

A poorly performing system can degrade the performance of all databases and applications deployed on that system. No amount of database, application, or SQL tuning can improve performance when a poorly implemented system is causing performance problems. Applications access databases and both are implemented on an overall system environment, as shown in [Figure 10.1](#). Therefore, a system problem can cause all databases and applications to perform poorly, just as a database problem can cause all applications that access that database to perform poorly.




Figure 10.1. The tuning boxes

A system problem can cause all databases and applications to perform poorly.

The system comprises the hardware and software required for the DBMS to operate and for applications to access databases using the DBMS. It is imperative that the DBA understand the system and operating environment where the database applications will be run. The DBA must be able to facilitate changes to any component of the system to tune the database environment. Of course, the DBA cannot be expected to be an expert in every aspect of the system, and therefore the DBA will need to work with other teams within the organization to initiate system changes.

The following sections provide introductory coverage of system-related performance and tuning tactics. They do not provide in-depth treatment of the subjects because that is not the primary focus of the book. However, we will

discuss system configuration issues, memory usage, database log configuration, the system catalog, and locking issues.

The Larger Environment

A DBMS operates within the context of a much larger environment that consists of other software and hardware components. Each of these components must be installed, configured, and managed effectively for the DBMS to function as required. The DBA needs to understand how the DBMS interacts with the server hardware, the operating system, and any other required software. Tuning and configuring these components and connections properly can have a dramatic impact on system performance.

A DBMS operates within the context of a larger environment.

Interaction with the Operating System

When the operating system experiences a performance problem, all of the software that runs on that operating system may experience performance problems. To help ensure an optimal operating system for database applications, the DBA should ask the following questions:

- Has a sufficient amount of memory been allocated for operating system tasks?
- Most operating systems have the capability of allocating a specific amount of disk space as a swap area. The swap area is used when the OS runs out of memory. Has a sufficient amount of disk space been allocated to the swap area?
- How were the database files allocated when the database was implemented? Interaction with the file system can cause some operating systems to create additional overhead. Changing the database files to use raw disk can eliminate OS and file system overhead. (Additional information on raw disk usage can be found in [Chapter 18, “Data and Storage Management.”](#))
- Some operating systems allow the administrator to set the priority of tasks that run under the auspices of the OS. Has each database-related task been assigned a priority? Is the priority appropriate for that specific task?
- Is the operating system at the version and release level recommended

by the DBMS vendor? Have any bug fixes been shipped for the OS that are applicable for the particular brand of database server you are running?

- Have the operating system configuration parameters been modified when installing the DBMS? If so, has sufficient testing been done to ensure that the parameters were modified correctly and do not impact any other processes that run on the database server?

Allied Agents

As discussed in previous chapters, the DBMS has to ally itself with many other software components to deliver service to the end user. Examples of allied agent software include

The DBMS has to ally itself with many other software components to deliver service to the end user.

- Transaction processors such as CICS and Tuxedo
- Application servers such as WebSphere or WebLogic
- Integration and ETL software such as SQL Server Integration Services (SSIS)
- Networking software such as TCP/IP and SNA
- Message queuing software such as MQSeries and MSMQ
- Web connectivity and development software such as ColdFusion and Rails
- IDEs and ORMs such as Visual Studio, Hibernate, Spring, and OpenAccess
- Programming languages such as Java, COBOL, and C

Each of these allied agents needs to be configured properly to interact with the DBMS, and it is the DBA's responsibility to understand the setup requirements. In larger shops the DBA might not perform the actual configuration—leaving it, instead, to more skilled professionals who specialize in administering and managing the software. However, in smaller shops DBAs may have to configure all of the software themselves.

Hardware Configuration

The DBMS runs on computer hardware. That hardware may be a large-scale

mainframe, an intermediate UNIX system, or a PC running Windows. Regardless of its scale, the hardware must be installed and set up properly for the DBMS to operate efficiently.

The hardware must be installed and set up properly for the DBMS to operate efficiently.

Again, here are some questions the DBA should ask to assure an optimal hardware environment for the database applications:

- Are the computer hardware and capacity appropriate for the DBMS environment? In other words, does the DBMS vendor recommend this particular hardware implementation?
- Is the computer firmware (e.g., ROM BIOS) up-to-date?
- Has a sufficient amount of memory been installed for all of the system software to be installed (OS, DBMS, and other allied agents)?
- Has an appropriate amount of disk storage space been allocated and configured for use by the DBMS?
- What type of disk storage is being used, and is it appropriate for large data volumes and high-speed database queries?
- Are all the network cables connected and functioning properly?
- Are all physical connections (e.g., cables, plugs, and board sockets) fully connected and operational?
- Is the hardware connected to an uninterruptible power supply?
- Is the hardware connected to a surge protection device?

Disk Storage and I/O

One of the biggest bottlenecks for database performance is the physical cost of performing I/O operations. Data resides on a disk, and a disk is a mechanical device. It requires machine parts that move in order to read encoded data from a spinning platter. This physical movement takes time, and anything that can be done to reduce I/O time can enhance performance.

A consideration for optimizing disk access is to use solid-state devices (SSDs). A solid-state device is actually computer memory that is configured to work like a disk drive. When data is read from a solid-state device, there is no physical component to the I/O operation—the data resides in memory and is transferred from memory to the DBMS and then to the requester.

A consideration for optimizing disk access is to use solid-state devices.

Consider placing database objects with high performance requirements on SSDs instead of physical disk drives, RAID devices, or storage area networks.

However, implementing SSDs has some potential problems. The first is cost. Solid-state devices will cost considerably more to acquire than traditional disk devices. The second potential problem is persistence. Some solid-state devices require a constant supply of power to prevent the data from being erased. In such cases, be sure that solid backup and recovery plans are implemented for database objects. Additionally, the longevity of flash-based solid-state devices is not up to par with traditional disk devices. For data that is modified frequently, keep in mind that flash-based SSDs can begin to degrade between 1 million and 5 millions writes. Be that as it may, software controllers can mitigate this issue, enabling flash SSDs to be viable for multiple decades before failure. Note that DRAM SSDs do not have this problem.

Consult [Table 10.1](#) for a summary of the characteristics of solid-state versus traditional disk-based storage.

Table 10.1. Solid State versus Traditional Disk

Performance or Implementation Consideration	Solid State	Traditional Disk
Spin-up time	Almost immediate as there are no mechanical parts	May take several seconds
Random access	0.1ms	Ranges from 5 to 10ms
Read latency	Low	High
Read performance	Consistent	Varying based on fragmentation
Defragmentation	Little benefit as read time is so rapid	Required after continued operation with data modification
Environmental consideration	Quiet with no moving parts and minimal susceptibility to impact	Mechanical parts can be loud and damage can result from impact
Longevity	Limited number of writes: 1M to 5M (if based on flash memory)	Mechanical failure but no limit on write/rewrite
Capacity	Up to 2TB (but more commonly smaller)	Up to 3TB (but more commonly 1TB to 2TB)
Power consumption	Typically half to one-third the power of a hard drive	12-18 watts for high-performance hard drives
Cost	Approximately \$2 per gigabyte	From 5 to 10 cents per gigabyte

Components of the DBMS

A DBMS is a very complex system requiring hundreds of thousands of lines of computer code. A DBMS is so complex that multiple programs are required to deliver the requisite data management functionality; each program interoperates with other programs to provide a database management system.

Each DBMS vendor breaks down DBMS functionality into different components. The DBA must study the makeup of the DBMS and gain an understanding of each component piece and how it integrates with the other components of the DBMS. For a high-level overview of the architecture of

the Oracle DBMS, refer to the sidebar “[The Architecture of Oracle.](#)”

The DBA must become an expert on the inner workings of the DBMS in order to ensure an optimized environment for database applications. A failure or problem in any single component of the DBMS can cause severe performance degradation for every application accessing the database.

The Architecture of Oracle

To effectively administer an Oracle environment, the DBA must understand the basic “architectural” blueprint of Oracle. Oracle is composed of five basic components that operate in an integrated manner to provide support for the client tasks: file structures, memory structures, processes, rollback segments, and redo logs.

An Oracle *instance* is the combination of all of the memory structures and background processes that are allocated when an Oracle database is started. Oracle users frequently confuse an Oracle *instance* with an Oracle *database*. An Oracle *database* has both physical structures (data files) and logical structures (table, index). The physical structure of an Oracle database is determined by the files created at the operating system level by the Oracle instance during database creation (e.g., controlfile, logfile) or by the DBA during normal operation (e.g., CREATE tablespace, CREATE controlfile). An Oracle database comprises three physical file structures:

- *Control files*—small administrative files that are used by the Oracle database.
- *Redo log files*—a record of changes made to data. The redo log files are used in recovery circumstances to ensure that no data is lost should a failure prevent changed data from being written to disk.
- *Database files*—files that contain the database information, including both system and application data.

An Oracle *parameter file* contains all configuration parameters for an Oracle instance. These parameters can be set to different values to tailor the Oracle instance to the system’s operating configuration as well as to the needs of the individual application(s) contained in the Oracle database.

Oracle *database files* contain the data associated with a particular

database. All of the files discussed are not absolutely required for normal database operations, but this configuration is highly recommended for a well-designed, efficient environment. Oracle database files can be grouped by the following categories:

- *Control files* record the physical structure of the Oracle database. The control file is read during every database start-up.
- *Data files* associated with Oracle tablespaces include system data files, application data files, default data files, temporary files, and rollback files.
- *Redo log files* (the database transaction log) record changes made to data.
- The *config.ora file*, associated with the Oracle client, specifies certain defaults, file, and directory names for the Oracle client.

Additionally, Oracle uses specific memory structures to perform DBMS-related tasks. These memory structures are contained in the main memory (RAM) of the computer running the Oracle instance. The basic memory structures for an Oracle instance are the system global area (SGA), the program global area (PGA), and the sort area.

The SGA is a group of shared memory structures allocated by the Oracle instance to contain data and control information. The SGA contains the data cache, redo log buffer, and shared pool for SQL statement parsing and processing.

The PGA is a work area for both user and background processes. Each process has its own PGA. The contents of the PGA will vary, depending on the type of process and Oracle configuration.

The Oracle sort area is a memory structure used to sort, order, and group data whenever a user process requests a sort to be performed.

Finally, we come to the Oracle processes, where most of the data management work is accomplished. Each process is composed of a series of tasks. Oracle has two general types of processes: user processes and Oracle processes. A *user process* is created to execute the program code of an application program. An *Oracle process* is called by another process to perform specific functions on behalf of the invoking process. Oracle processes can be further broken down into server processes and background processes. *Server processes* communicate with user

processes, acting as a “relay” between the user process and SGA information. *Background processes* perform designated data management and processing functions for the Oracle instance.

Let’s examine the functionality of each of the Oracle background processes.

The *process monitor* (PMON) background process performs “cleanup duties” when a user process fails with an error condition. PMON cleans up the cache, releases locks, and performs other miscellaneous tasks.

The *system monitor* (SMON) background process provides instance recovery during start-up. SMON also cleans up temporary segments that are no longer in use, compresses free space into contiguous extents, and, in a parallel server environment, provides instance recovery for a failed CPU.

The *database writer* (DBWR) process writes data from the data cache contained in memory out to the physical disk files.

The *log writer* (LGWR) manages the redo log buffer. If an Oracle instance takes more checkpoints than the LGWR process can handle, checkpoint duties may be turned over to the checkpoint (CKPT) background process. The archiver (ARCH) process performs log archival by copying online redo log files to auxiliary storage devices as logs become filled up.

The *recover* (RECO) background process automatically resolves failures involving distributed transactions.

Finally, Oracle deploys *user* and *server* processes. The user process links the client application to the Oracle instance. The server process parses and executes SQL statements and performs all tasks required to send results back to the requesting application.

This overview of Oracle’s architecture is necessarily brief. It is intended to communicate just how specialized the components of a DBMS must be. An Oracle DBA will need to understand all of these components, how they interact, and how they can be tuned to optimize the performance of the Oracle system and applications.

DBMS Installation and Configuration Issues

Every database management system provides parameters that allow the DBA to configure various aspects of the database environment. Configuration is accomplished in a variety of ways, depending on the DBMS. Some popular configuration methods include executing system procedures to set and reset values, editing files with parameter settings, issuing commands at a DBMS prompt, and assembling parameter specifications using an option within the DBMS. Regardless of the manner of configuring the DBMS, the DBA will be required to specify various parameters to the DBMS that affect the way the DBMS operates.

Most DBMS software ships with default values supplied by the vendor. However, default values are usually not sufficient to support the development of robust production applications. This section will discuss some of the common configuration parameters and how to tune them.

Default values are usually not sufficient to support the development of robust production applications.

Types of Configuration

The DBMS can be configured when the DBMS is installed or after the DBMS is operational. During installation, the DBA or SA installing the DBMS will have the option to change the configuration parameters or to allow the parameters to default. However, defaults are almost always wrong. It is almost always better to set the parameters based on knowledge of the data and applications that will be using the DBMS.

Once the DBMS is installed and operational, the DBMS will provide a method of changing the configuration parameters. Depending on the DBMS, parameters may be changed dynamically, nondynamically, or both.

Depending on the DBMS, parameters may be changed dynamically, non-dynamically, or both.

Dynamic parameters can take effect without the need to restart the database server. For example, executing the RECONFIGURE command in Microsoft SQL Server causes dynamic parameters to immediately take effect. After the appropriate commands are issued, the parameter value changes and the DBMS will behave accordingly.

In order for nondynamic parameters to take effect, the DBMS must be shut

down and restarted. Of course, the parameter value must be changed—usually in the same or a similar manner as a dynamic parameter. However, a nondynamic parameter will not take effect until the DBMS is restarted.

Memory Usage

Relational databases love memory. The single biggest system performance tuning task that a DBA will face is configuring RDBMS memory usage. The DBMS uses random access memory to cache data and other resources required by the DBMS. This is done because reading data from memory is much less costly than reading the data from disk. So, as a rule of thumb, the more memory you can provide to the DBMS, the better performance will be. Of course, the DBMS has to be configured properly to use the memory efficiently.

Relational databases love memory.

The *American Heritage Dictionary* defines [cache](#) as “a place for concealing and safekeeping valuables.” This is a good starting point for understanding what a cache means to a relational database. The “place” used by the DBMS is memory (as opposed to disk). The “valuables” are data pages, query plans, and other database resources. A typical DBMS will use a cache to keep resources in memory longer. The longer a resource is cached in a memory location, the better the chance for subsequent requests for the resource to avoid incurring costly I/O operations, as shown in [Figure 10.2](#).

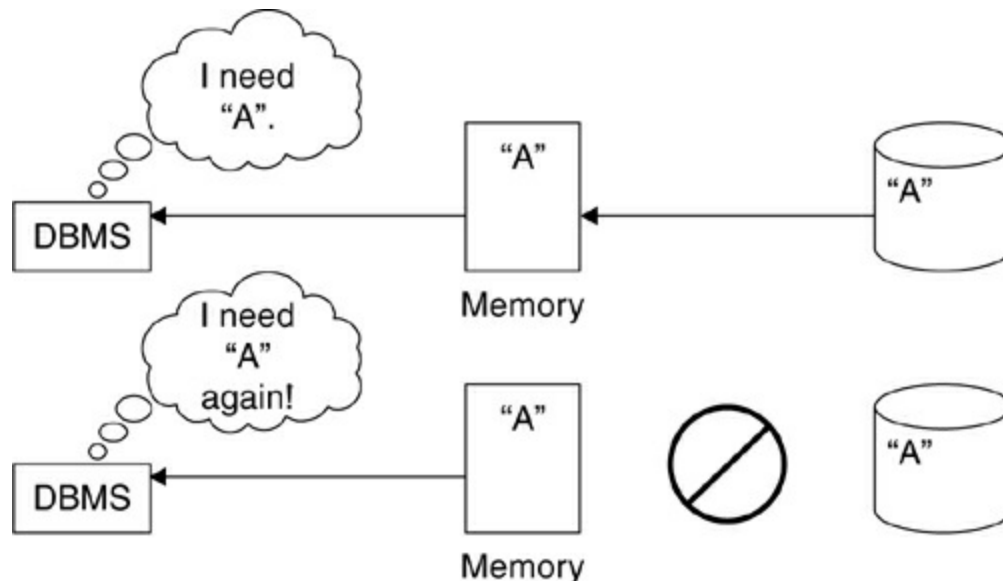


Figure 10.2. The value of caching resources in memory

There are multiple caches or buffers used by DBMS products to reduce the cost of I/O. Each DBMS uses different terminology but generally caches the same resources.

A *data cache* is used to avoid I/O operations when actual data is being read from the database. Accessing data in memory is substantially faster than accessing data on disk. Therefore, all data access from a relational database goes through a caching area (see [Figure 10.3](#)). Access to memory is typically measured in microseconds, while access to disk I/O is usually measured in milliseconds. When a program needs a row from a table, the DBMS retrieves the page from disk and stores the page in the data cache. Basically, the DBMS uses the cache as a staging area. If the row changes, the change is written to the page in the data cache. Eventually, the DBMS will write the page in the data cache back to disk. When data needed by an application is on a page already in the data cache, the application will not have to wait for the page to be retrieved from disk. Depending on the DBMS, this memory structure also may be called a buffer pool.

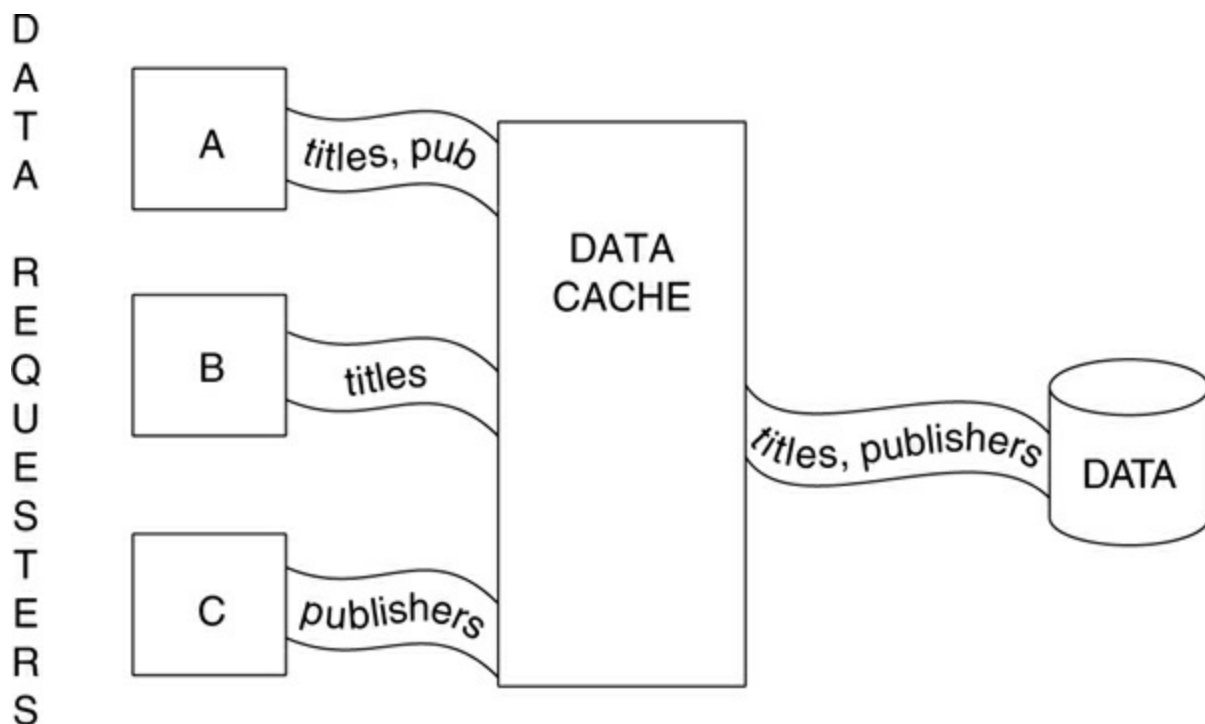


Figure 10.3. Data cache (or buffer pool)

A data cache is used to avoid I/O operations when actual data

is being read from the database.

A *procedure cache* stores SQL and program-related structures. Before SQL statements can be issued to retrieve or modify data, the statement must first be optimized by the DBMS. The optimization process creates an internal structure representing the access path that will be used by the DBMS to read the data. The DBMS can store these access paths in the procedure cache and reuse them each time the program or SQL statement is run. This optimizes application performance because the optimization process need not be performed every time the SQL is run. Instead, optimization occurs the first time the SQL is issued, and subsequent executions retrieve the access path from the procedure cache. Each DBMS provides similar functionality, although with different names and differing features.

Another memory structure commonly deployed within a DBMS is a *sort cache*. The sort cache is used instead of temporary disk storage to store intermediate sort results in memory. The more sorting functionality that can be performed in memory, the better a sort will perform. Many relational database operations require sorts, for example, grouping, ordering, UNION operations, and certain types of joins.

The sort cache is used instead of temporary disk storage to store intermediate sort results in memory.

The DBMS may also use other *internal structure caches*. The implementation of each DBMS is unique. To accomplish relational operations, the DBMS may need to create internal structures that are not necessarily visible to the end user. However, DBAs, and sometimes programmers, will need to know about the internal structures. One example is the internal DBD (database descriptor) structure used by DB2 to manage databases. The DBD is never externalized to users, but every time an application accesses any object within a database, DB2 must load the DBD into a memory area known as the EDM pool. DB2 uses EDM pools to cache dynamic SQL access paths and other internal structures, as well. DB2 DBAs need to allocate sufficient memory and monitor the efficiency of the EDM pools as processing requirements and usage patterns change.

The DBMS also may buffer log records to a separate *database log cache*. Furthermore, the DBMS may implement two log caches, one for log writes

and one for log reads. The database log keeps a record of all changes made to the database. The log write cache is used to speed up database modifications. The changed data is written to the log cache, and over time the cache is written asynchronously to disk. When log writes are buffered in this way, the database log becomes less of a bottleneck to system and application performance. The log read cache is used for ROLLBACK and RECOVER operations. A rollback or a recovery needs to access the log to undo or reapply database changes. As the log records are requested, they will be buffered in memory in the log read cache.

The DBMS also may buffer log records to a separate database log cache.

Additional Areas of Memory Consumption

In addition to the various caches and buffer pools used by relational database systems, memory is required for other purposes. Generally, the DBMS installation or configuration routines allow the DBA to allocate and tune the memory consumption of the DBMS. Some of the more common areas of DBMS memory consumption include

- *User connections.* Each concurrent user connection to the DBMS, regardless of the type of client connection, requires memory for the DBMS to maintain and manage the connection.
- *Devices.* The devices used by databases may require system memory to maintain and use.
- *Open databases.* Most DBMSs provide a parameter to specify the maximum number of databases that can be open at any one time. Each open database requires DBMS memory.
- *Open objects.* Another parameter may exist to identify the maximum number of database objects that can be open at any one time, including tables, indexes, and any other database object in use. Each open database object requires memory.
- *Locks.* Each concurrently held lock will require memory. The DBMS should provide a configuration parameter for the number of concurrent locks that can be held at one time.
- *Caches.* The various caches are discussed in the previous section.

How Much Memory Is Enough?

So, if relational databases love memory, just how much memory should be allocated? This is a difficult (if not impossible) question to answer by using generalities. The tempting answer is “Enough to get the job done,” but that does not help the DBA who has to allocate the right amount of memory in the right place.

Every DBMS uses memory, but in different amounts and for different things. The best approach is to search your vendor’s DBMS manual to determine how much memory is required for each resource. Then you can estimate the usage requirements for each resource and calculate the approximate amount of memory required by the DBMS.

Every DBMS uses memory, but in different amounts and for different things.

For example, SQL Server requires about 75 bytes of memory per lock. To configure the amount of memory required for locking, the DBA will need to estimate the total number of concurrent transactions and the average number of locks per transaction. After this has been done, the calculation for memory required by the DBMS for locking is simple:

$$\text{Concurrent locks} = [\# \text{ concurrent transactions}] \times [\# \text{ locks per transaction}] \times 75$$

The result of this calculation provides an estimate for the memory (in bytes) used for locking in a SQL Server system. Of course, the number of concurrent transactions and the number of locks per transaction may be difficult to ascertain for your environment. You may be able to substitute the number of concurrent transactions with the DBMS configuration value used for total number of user connections. If the number of locks per transaction is unavailable, you can examine a few typical programs to come up with an estimate. This process is then repeated for each database resource that consumes memory. Of course, the data cache should be treated separately and is covered in the next section.

After repeating this process for each resource, you will arrive at the amount of memory to install for the database server. It is a good practice to leave some breathing room to support the addition of unanticipated application programs, DBMS and operating system upgrades requiring additional memory, and other unforeseen items. This means installing a little bit more memory than you calculate. However, do not leave too much

breathing room. The memory was installed to be used, not conserved.

It is a good practice to leave some breathing room.

Data Cache Details

It is important for the DBA to understand the basics of data cache operation. At any given point in time, the data cache will consist of three different types of pages:

- *In-use pages*—pages that are currently being read and updated by the DBMS. These pages are not available to be used for other database processing.
- *Updated pages*—pages where data has been modified in some way, but the data has not yet been written to disk. These pages are not available to be used for other database processing.
- *Available pages*—pages not being used. These pages are available to be used for other database processing. New data can be written to available pages in the data cache.

The performance of the data cache depends on how efficiently memory has been allocated for its use. For example, if unavailable pages dominate the data cache, the DBMS may trigger synchronous writes to increase the space in the data cache. Synchronous writes can slow down database processing because each write request has to wait for the data to be actually physically written to disk. Depending on the type of processing the database applications are using, the DBA may be able to tune data caching parameters and sizes to enable more efficient buffering of data.

Monitoring and Tuning the Data Cache

Ensuring the cache is the proper size is one of the most critical aspects of having an efficient data cache. A data cache that is too large wastes memory and can cause pages in the cache to be moved out to auxiliary storage. A data cache that is too small forces frequent writes to disk and, in the most severe cases, results in swapping the data cache pages back and forth from disk.

Ensuring the cache is the proper size is critical.

The complexity of tuning the data cache varies by DBMS and depends on the configuration and tuning parameters that are available for the data cache.

Some DBMSs, such as DB2, provide multiple buffer pools that can be configured and tuned independently with multiple parameters. Others, such as SQL Server, are more basic, with a data cache per database. But regardless of the DBMS, the DBA should monitor the read efficiency of each data cache or buffer pool.

The read efficiency of the data cache is a percentage that tracks how well the cache is performing its primary duty—to avoid physical disk I/O operations. The read efficiency of each data cache is calculated as the number of actual I/Os performed subtracted from the total number of data requests, then divided by the total number of data requests, or

$$\text{ReadEfficiency} = \frac{(\# \text{ database I/O requests}) - (\# \text{ physical I/Os})}{(\# \text{ database I/O requests})}$$

In other words, read efficiency shows the percentage of times a data page is found in the data cache (or buffer pool). The higher this percentage is, the more efficient the buffer pool is. When data pages can be found in buffers in memory without requiring a physical I/O, performance will be enhanced.

The actual numbers for I/O requests and actual physical I/O operations can be found by examining DBMS trace records or by using a database performance monitor. Be sure to examine all I/O requests, those made synchronously as well as those made asynchronously. Many DBMSs are capable of anticipating I/O requirements and scheduling reads of multiple blocks prior to their being requested. For example, DB2 deploys sequential prefetch to read multiple pages into the buffer pool for sequential reads that access large amounts of data.

Depending on the DBMS, the DBA may need to turn on traces that externalize instrumentation details of the DBMS. Each DBMS has its own set of instrumentation information that can be examined. Also, depending on the type of performance monitor, traces may not need to be started by the DBA because the monitor may start and stop traces as required. The monitor may also use other means to capture the performance information from the DBMS.

As a rule of thumb, an 80 percent or better read efficiency is good for a data cache. Of course, the read efficiency value will depend on the type of processing. Many sequential processes can cause the data cache to be overrun and the efficiency to drop. Furthermore, systems with many processes that

access data only weekly or monthly may have lower read efficiencies because less data is frequently reused. The DBA should know the type of database processing that occurs for each data cache and gauge the read efficiency of that data cache in light of the processing that occurs.

An 80 percent or better read efficiency is good for a data cache.

When read efficiency is consistently below 80 percent, consider tuning by increasing the size of the data cache, modifying the number or type of concurrent processes using the data cache, or reducing the number of tables and indexes assigned to the data cache. Making such changes may impact availability because the DBMS must be stopped and restarted to register the change. Each DBMS has different requirements for making data cache changes.

Depending on the DBMS, the DBA may be able to configure the data cache better by changing the amount of cache reserved for sequential and parallel operations. Large table scans can quickly monopolize a data cache. Reserving only a subset of the entire cache for sequential scans may improve the overall performance of the data cache. Of course, the reverse may be true if the majority of operations that use the data cache are sequential in nature.

Large table scans can quickly monopolize a data cache.

Additional possibilities for tuning the data cache—offered by some DBMS vendors or as add-on products—include creating an additional cache to back up the data cache, pegging specific pages in memory, and automatically growing the size of the data cache as throughput increases.

Monitoring and Tuning the Procedure Cache

The DBA must monitor the effectiveness of the procedure cache to help improve the efficiency of database applications and queries. Although the procedure cache differs dramatically from DBMS to DBMS, the general idea is the same: to keep optimized SQL structures in memory so they can be reused by subsequent tasks instead of being reformulated or reread from disk.

To ensure optimal performance, the procedure cache must be sized properly to accommodate all the SQL that may be run concurrently. DBAs need information about the application programs that will be run in order to

size the procedure cache effectively. The read-efficiency calculation we used for gauging the effectiveness of the data cache can be used for the procedure cache also. In the case of the procedure cache, the read efficiency calculates how often the DBMS needs to reoptimize SQL. Procedure cache read efficiency usually ranges from 60 to 80 percent. Of course, this percentage range will differ with the type of DBMS and the type of caching performed, as well as with the type of applications and the number of times the same SQL and programs are run. Refer to the sidebar “[DB2’s EDM Pools](#)” for an example of a procedure cache.

The procedure cache must be sized properly to accommodate all the SQL that may be run concurrently.

DB2’s EDM Pools

IBM’s DB2 for z/OS caches SQL in a memory area known as the *environmental descriptor manager pools*, or EDM pools for short. The EDM pools contain active and skeleton structures for application plans and packages, as well as optimized dynamic SQL plans, database descriptors, and program authorization lists. The EDM pools function as a sort of catchall memory area for resources that are used by application programs and SQL statements as they are being run.

If your EDM pools are too small, DB2 will incur increased I/O activity in the DB2 directory, a specialized system catalog for internal DB2 structures. An additional symptom that may occur is increased response times as the required structures are read from disk and loaded into memory in the EDM pools.

“Open” Database Objects

As database applications process, the DBMS will need to open database objects for access and maintain the open objects. The DBMS will probably have a configuration option for setting this number. However, these numbers are difficult to arrive at originally. As time progresses, the database system can be monitored and better values can be provided.

To begin with, it is not uncommon simply to specify a value such that all production databases and objects can be open. However, each open database object consumes memory. Therefore, it is a good idea to temper this value,

taking into account the size of the database implementation, the type of application processing, and the amount of memory available.

Each open database object consumes memory.

Database Logs

The database log is another database system configuration option that can impact performance. The database log, sometimes referred to as the *transaction log*, is a fundamental component of a database management system. All changes to application data in the database are recorded serially in the database log (see [Figure 10.4](#)). Using this information, the DBMS can track which transaction made which changes to the database. Furthermore, ROLLBACK and RECOVER operations use the database log to reset the database to a particular point in time.

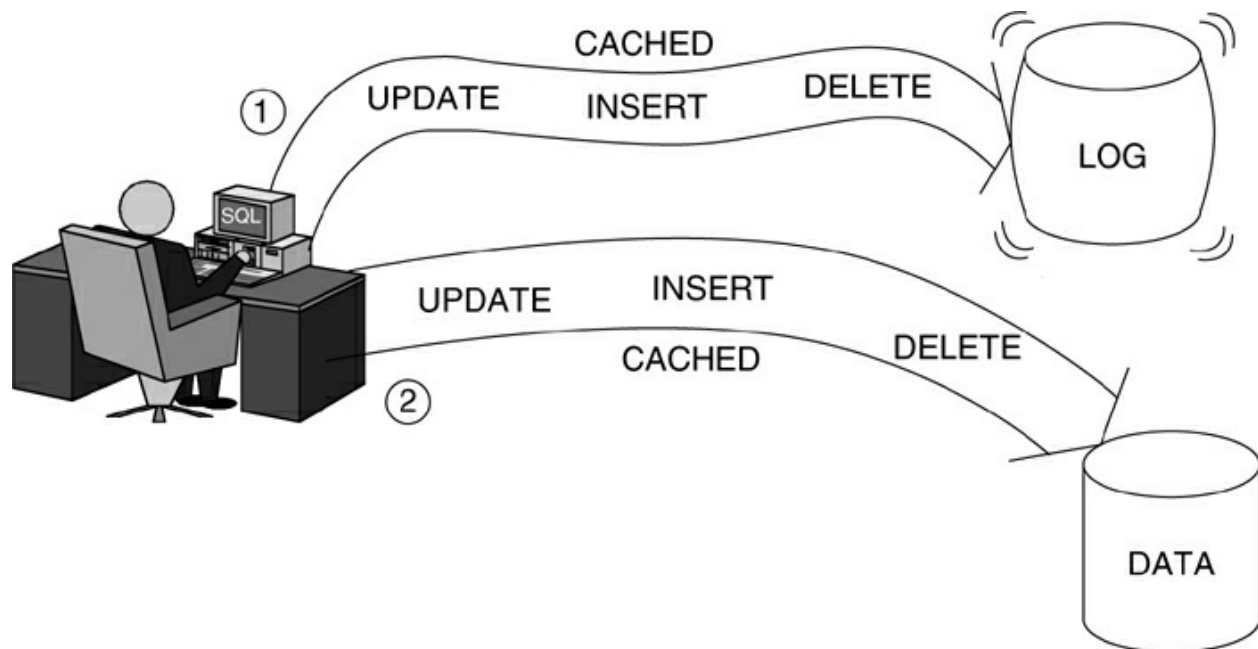


Figure 10.4. Logging of database transactions

All changes to application data in the database are recorded serially in the database log.

The manner in which the database log is created depends on the DBMS. Some DBMSs specify the log at the database system level; others define a database log for each database that is created within the database system. Some DBMSs provide a parameter to enable and disable logging. In general,

avoid disabling database logging for any database or database system where the data is valuable. In other words, consider turning off database logging only for test databases.

Depending on the DBMS, the database log may consist of several files. For example, Oracle uses a transaction log and rollback segments to accomplish all of the functionality described in this section.

During normal database application processing, SQL INSERTs, UPDATEs, and DELETEs will be issued to modify data in the database. As these database modifications are made, the transaction log for the database will grow. Since each database change is logged, the DBA will need to monitor actively the size of the transaction log files. Since data is constantly changing, the log will be continually growing.

The database transaction log is a write-ahead log.

The database transaction log is a *write-ahead log*. This means that changes are made to the transaction log before they are made to the data in the database tables. When the database modification has been fully recorded on the log, recovery of the transaction is guaranteed.

Typically, the DBMS takes a *system checkpoint* to guarantee that all log records and all modified database pages are written safely to disk. The frequency of database system checkpoints can be set up by the DBA using database configuration parameters. Checkpoint frequency is usually set as either a predetermined time interval or a preset number of log records written.

Generally, the following type of information is recorded on the database log:

- Beginning and ending time of each transaction
- Actual changes made to the data and enough information (using “before” and “after” images of the data) to undo the modifications made during each transaction
- Allocation and deallocation of database pages
- Actual COMMIT or ROLLBACK of each transaction

Using this information, the DBMS can accomplish data integrity operations to ensure consistent data. The transaction log is used when the DBMS is restarted, when transactions are rolled back, and to restore a

database to a prior state. Let's examine each of these scenarios.

When the DBMS is restarted, each database goes through a recovery process. During restart processing, the DBMS will check to determine which transactions must be rolled forward. This occurs for transactions where it is unknown if all the modifications were actually written from the cache to disk. A checkpoint forces all modified pages to disk. Therefore, it represents the point at which the start-up recovery must start to roll transactions forward. Because all pages modified before the checkpoint are guaranteed to be written accurately to disk, there is no need to roll forward anything done before the checkpoint.

When a transaction is rolled back, the DBMS copies "before" images to the database of every modification made since the transaction began.

During a recovery scenario, the DBA can use the transaction log to restore a database. First, a backup copy of the database is restored, and then subsequent transaction log backups can be restored. This causes a roll forward of the transaction log. During a roll forward, the DBMS will copy to the database "after" images of each modification. Using the logged data, the DBMS ensures that each modification is applied in the same order that it originally occurred.

During a recovery scenario, the DBA can use the transaction log to restore a database.

You can see that the transaction log is a useful item to have around in case of database or transaction errors, and to ensure data integrity.

Database Log Configuration Considerations

The configuration of the database log can be a complicated task. Depending on the DBMS, the DBA may need to make multiple configuration decisions to set up the database transaction log, such as defining input buffers and output buffers, setting up log off-loading, and defining the actual log files.

Defining output buffers for the database log can optimize log write operations. Writing log records to memory instead of directly to disk creates more efficient database processing. The DBMS can write log records asynchronously from the output buffer to the physical log file later. When logging is implemented in this fashion, database processing does not need to wait for synchronous log writes to occur to disk.

Defining input buffers for the database log can optimize operations—such as ROLLBACK and RECOVER—that read the database log.

It is a good idea to set up dual logs.

When configuring the database log, it is a good idea to set up dual logs. With dual logs, the DBMS will log changes to two separate and independent log files. Implementing dual logging provides a redundant log for use in case one of the logs fails. A log can fail for many reasons, including device failure, media problems, or simple carelessness. When setting up dual logging, be sure to define each log on separate devices and on separate controllers to minimize the possibility of both logs failing at the same time.

Another log configuration detail is deciding how database logs are to be handled when they fill up. Once again, the implementation details depend on the DBMS. Some DBMSs provide for automatic log off-loading. [Log off-loading](#) is the process of archiving an active log to an archival log and switching log writes to a new active log.

If the DBMS performs automatic log off-loading, you can improve performance by off-loading to an archive log on disk instead of tape. Archiving to disk enables the log off-loading process to run faster, and backout and recovery processes can also run faster because the log records will be on disk—meaning not just faster I/O but no wait for tapes to be mounted. The DBA can use a storage management system to automatically migrate archive logs to tape after a predetermined amount of time.

Some DBMSs require the database log to be explicitly backed up. In this case, you will need to implement periodic backups of the transaction log. Typically, the DBMS provides a specific backup command to create a transaction log backup. When the DBMS finishes backing up the transaction log, it truncates the inactive portion of the transaction log to free up space. Truncation of the log enables the DBMS to reuse space.

Of course, as a DBA you will need to learn how each DBMS you use handles backing up the database log. DB2 automatically archives the transaction log; Microsoft SQL Server requires the DBA to back up transaction log files. In fact, Microsoft SQL Server provides some interesting features for handling the transaction log; the sidebar “[Microsoft SQL Server Log Options](#)” discusses the parameters that can impact the operation of the database log.

Are All Database Operations Logged?

Depending on the DBMS, certain situations and commands may not be logged. To avoid “out of space” conditions caused by rapid growth in transaction log files, the DBMS may turn off logging. For example, some DDL operations and database utility executions might not be logged. Beware of these situations and plan accordingly for your recovery needs.

Certain situations and commands may not be logged.

Microsoft SQL Server Log Options

Microsoft SQL Server provides a system configuration parameter that affects the behavior of database logging. The TRUNC LOG ON CHKPT option can be changed at the database level. Use the system procedure named SP_DBOPTION to change the configuration settings for a database. For example:

[Click here to view code image](#)

```
EXEC SP_DBOPTION 'pubs', 'trunc. log on chkpt.',  
'false'
```

Issuing this command causes the TRUNC LOG ON CHKPT option to be set to FALSE for the pubs database. To see a list of all current database options set for a database, simply issue the system procedure without additional parameters, for example:

```
EXEC SP_DBOPTION pubs
```

The TRUNC LOG ON CHKPT option is potentially quite dangerous. When this option is set to TRUE, every checkpoint operation will cause the database log to be truncated—that is, the database log will be emptied and reset, causing all of the logged changes to be lost. A database cannot be recovered using a log that is truncated on checkpoints.

You might consider setting the TRUNC LOG ON CHKPT option to TRUE for test databases during your application development cycle, but not for mission-critical production databases. The DBA needs to be aware of these options and ensure that the appropriate settings are made, based on the needs of the users of the data and the applications that

access that data.

You can also control log truncation using the SQL Server recovery models. A *recovery model* is a database property that controls how transactions are logged, whether the transaction log requires (and allows) backing up, and what kinds of restore operations are available. Three recovery models exist: simple, full, and bulk logged. Enabling simple recovery mode causes log space to be reclaimed automatically, whereas full and bulk-logged recovery modes do not automatically truncate the transaction log.

During some large operations, such as CREATE INDEX, the DBMS will probably not log every new page. Instead, the DBMS will record enough information in the database log to determine that a CREATE INDEX happened, so that it can either be recreated during a roll forward or removed during a rollback.

Additionally, some DBMSs provide configuration options at the database level to turn off logging for some types of processing. For example, in Microsoft SQL Server, when the SELECT INTO/BULKCOPY database option is set to TRUE,¹ the following operations will not be recorded in the database transaction log: bulk load operations, TRUNCATE TABLE, BULK INSERT, and SELECT INTO statements. These operations usually cause a large volume of data to be changed in the database. Logging can slow down these processes, so logging can optionally be disabled for these operations only. However, because these operations are not recorded in the transaction log, SQL Server cannot use the RESTORE operation on the transaction log to recover them. If the DBMS has no knowledge of the operations in the log, it cannot recover the data.

Of course, SQL Server is not the only DBMS that disables logging for certain operations. Another example is the DB2 REORG utility. A parameter is provided to disable or enable logging during the reorganization process. DB2 also provides for the option to create nonlogged tablespaces, for which no database logging will occur. Without logging, these objects cannot be recovered using the log, so care should be taken to deploy nonlogged objects only under specific scenarios (for example, to implement a queue that is filled and emptied during application processing).

As a DBA, you will need to learn how and when to turn off logging using

your DBMS. Moreover, keep in mind that whenever logging is turned off, the data must be backed up both before and after the nonlogged process to ensure point-in-time recoverability.

Locking and Contention

Concurrency operations such as *deadlock detection* and *lock manager settings* can greatly impact database system performance. Database processing relies on locking to ensure consistent data based on user requirements and to avoid losing data during updates. You must balance the need for concurrency with the need for performance. If at all possible, seek to minimize the following situations:

Balance the need for concurrency with the need for performance.

- *Lock suspensions* occur when an application process requests a lock that is already held by another application process and cannot be shared. The suspended process temporarily stops running until the requested lock becomes available.
- *Time-outs* occur when an application process is terminated because it has been suspended for longer than a preset interval. This interval can usually be set by using a configuration parameter.
- *Deadlocks* occur when two or more application processes hold locks on resources that the others need and without which they cannot proceed. The deadlock detection cycle—that is, the time interval between checking for deadlocks—can also be set by using a configuration parameter.

When accessing a relational database, the locking process can be quite complex. It depends on the type of processing, the lock size specified when the table was created, the isolation level of the program or SQL statement, the method of data access, and the DBMS configuration parameters. To tune database locking requires a combination of system, database, and application tuning.

The System Catalog

The physical location and setup of the system catalog will have an impact on system performance. The DBA must decide where it will be installed, on

what type of disk, and how much space to allocate. These decisions typically are made at installation time.

Place the system catalog on a separate disk device.

As a rule of thumb, place the system catalog on a separate disk device so that it can be managed and tuned independently from other application data. If possible, consider completely dedicating a disk volume or two to the system catalog. Consider placing the indexes and tables on separate disk volumes. In addition, if the DBMS does not already provide a separate data cache for the system catalog, consider isolating the system catalog into its own dedicated data cache. Doing this makes it easier to track the efficiency of system I/O versus application I/O.

A Catalog by Any Other Name . . .

A relational DBMS requires a system catalog, but it may be called something different depending upon the particular DBMS you are using:

- In Oracle it is referred to as the Data Dictionary.
- In SQL Server it is referred to as the System Catalog.
- In DB2 it is referred to as either the DB2 Catalog or System Catalog.
- In MySQL it is referred to as the Information Schema.

When changes need to be made to the system catalog database, utilities such as REORG, COPY, and RECOVER or file system commands need to be used. Changes may need to be made to increase the size of the system catalog, to add a new index, or to migrate to a new release of the DBMS. Usually a migration utility is provided to make system catalog changes.

DBAs should use the system catalog to actively manage their database environment. It is a good practice to actively monitor the database objects in the system catalog and delete obsolete objects. For example, in DB2 for z/OS, IBM delivers sample tables with every new version. The database object names have the DB2 version number embedded within them. Every new release causes new sample tables to be created. The DBA should delete the old sample tables when new ones are installed. This advice applies to all unused database objects, not just to sample tables. If a tablespace exists that is no longer used, it is consuming valuable resources (disk space, system

catalog space, etc.) that can be freed up when the database object is dropped.

Although you can directly query the system catalog tables using SQL, many DBMS products provide catalog views to simplify writing such queries.

Other Configuration Options

Every DBMS will have many configuration and tuning options that are particular to that DBMS. The DBA needs to become an expert in the options that are available and understand the impact of each permissible setting.

Some example configuration options that may be encountered include the following:

- *Nested trigger calls.* Some DBMSs can enable and disable nested trigger calls. A *nested trigger call* is when one trigger causes another trigger to fire. Some DBMSs may provide additional control over trigger nesting by providing a maximum value for it. By setting this value, the DBA can control how many levels of nested trigger calls are allowable. Control over triggers can have an enormous impact on performance. For example, if an application hits the maximum number of nested triggers, all of the changes caused by all previous triggers need to be rolled back—potentially causing considerable performance degradation.
- *Security options.* The functionality of security and authorization can be controlled by DBMS configuration options. Some DBMSs allow database security to be turned over to external security and control software.
- *Identity values.* The identity property can be assigned to a column such that the DBMS automatically assigns numerically sequential values when data is inserted in the table. The DBMS can allow the configuration of the pool size from which identity values are obtained.
- *Distributed database.* To configure a distributed database implementation, the DBMS most likely will provide options for connecting databases at various locations.

General Advice

When configuring your database environment, avoid defaults. Most configuration options will default to a predefined value if no specific value is

assigned. The major problem with using default values is that they almost never are the best choice for your particular environment. It is always a good idea to specify the value of every configuration parameter *even if the default value is the one you wish to choose*.

Avoid default configuration options.

Finally, beware of configuration options that change the behavior of the DBMS. Simply setting a single configuration parameter to the wrong value can cause a lot of damage. Consider the following parameters and their potential impact:

- Oracle provides an optimization-mode parameter to choose between cost-based and rule-based SQL optimization. These two methods can create dramatically different SQL access paths with dramatically different performance.²
- Sybase provides a parameter called ALLOW UPDATES that controls whether the system catalog tables can be modified using SQL. When this option is turned on, the system catalog can be changed quite easily. This is to be avoided except under the guidance of Sybase technical support. System catalog changes should be driven by the DBMS itself, not by user-issued SQL.
- DB2 provides several parameters that control the behavior of database operations such as parallel queries, data sharing, and dynamic SQL caching.

System Monitoring

A DBMS environment should be continually monitored for performance degradation and problems. Some DBMSs provide built-in monitors to accomplish rudimentary monitoring. Additional performance monitoring solutions such as Oracle Enterprise Manager, Foglight from Quest Software, or OMEGAMON from IBM Corporation can be deployed to provide more robust monitoring functionality, such as dynamic alerting and proactive management of performance problems. Refer to the sidebar “[Sybase SP_MONITOR](#)” for information on the Sybase DBMS performance monitor.

Sybase SP_MONITOR

The SP_MONITOR system procedure is a rudimentary tool supplied by Sybase with its DBMS. When the procedure is run, it displays DBMS activity in terms of the following performance-oriented numbers:

- **last run:** date and time that SP_MONITOR was last executed
- **current run:** date and time that SP_MONITOR was run to produce this report
- **seconds:** total number of seconds since the last SP_MONITOR was run
- **cpu busy:** CPU time, in milliseconds, that the server CPU spent performing Sybase work
- **io busy:** CPU time, in milliseconds, that the server CPU spent performing Sybase I/O operations
- **idle:** CPU time, in milliseconds, that Sybase has been idle
- **packets received:** number of input packets read by Sybase
- **packets sent:** number of output packets sent by Sybase
- **packet errors:** number of errors detected by Sybase while reading and writing packets
- **total read:** number of disk reads performed by Sybase
- **total write:** number of disk writes performed by Sybase
- **total errors:** total number of errors detected by Sybase while reading or writing
- **connections:** number of logins or attempted logins to Sybase

The numbers are cumulative since the last time SP_MONITOR was executed.

Performance monitors are available for every aspect of the system environment, not just the DBMS, so you may be able to monitor the operating system, network, transaction server, and any other system middleware for performance problems. A DBA must be able to operate and understand the output of the monitoring solutions available. The performance monitor acts as a window into the efficiency (or lack thereof) of the database system.

The performance monitor acts as a window into the efficiency of the database system.

Summary

If the environment where the DBMS must operate is not performing efficiently, it is impossible for the DBMS, and indeed any database access, to perform efficiently. The DBA needs to understand every DBMS configuration value and the impact it can have on the overall performance of the system. Furthermore, the DBA must control the integration of the DBMS with the hardware on which the DBMS runs and any allied agent software.

Review

1. What is the single most important configurable component of a relational database management system?
2. How does the data cache (or buffer pool) improve the performance of database processing?
3. Describe how to determine the read efficiency of the data cache.
4. What components of the DBMS require system memory to control?
5. What three concurrency problems can negatively impact performance?
6. Why are certain database operations not logged?
7. A DBMS is composed of multiple programs interacting with one another: true or false?
8. What are the three possible statuses of a page in the data cache?
9. What benefits can accrue by caching optimized SQL in memory?
10. What type of information is recorded on the database transaction log?

Bonus Question

When a database operation is not logged, what precautions should the DBA take before and after the nonlogged operation? Why?

Suggested Reading

Bruzzese, J. Peter, Ronald Barrett, and Wayne Dipchan. *Windows Server 2008: How To*. Indianapolis, IN: SAMS (2010). ISBN 978-0-672-33075-9.

Hassell, Jonathan. *Learning Windows Server 2003*. 2nd ed. Sepastopol, CA: O'Reilly (2006). ISBN 0-596-10123-6.

- Henderson, Ken. *The Guru's Guide to SQL Server Architecture and Internals*. Boston, MA: Addison-Wesley (2004). ISBN 0-201-70047-6.
- Johnson, Robert H. *MVS Concepts and Facilities*. New York, NY: McGraw-Hill (1989). ISBN 0-07-032673-8.
- Kirkwood, John. *Sybase Architecture and Administration*. New York, NY: Ellis Horwood (1993). ISBN 0-13-100330-5.
- Limoncelli, Thomas A., et al. *The Practice of System and Network Administration*. Boston, MA: Addison-Wesley (2007). ISBN 978-0-321-49266-1.
- Nemeth, Evi, et al. *UNIX and Linux System Administration Handbook*. 4th ed. Upper Saddle River, NJ: Prentice Hall (2010). ISBN 978-0-13-148005-6.
- Mullins, Craig S. *DB2 Developer's Guide*. 6th ed. Boston, MA: IBM Press (2012). ISBN 978-0-13-283642-5.
- Reese, George, et al. *Managing & Using MySQL*. 2nd ed. Sebastopol, CA: O'Reilly (2002). ISBN 0-596-00211-4.
- Reiss, Levi, and Joseph Radin. *Unix System Administration Guide*. Berkeley, CA: Osborne/McGraw-Hill (1993). ISBN 0-07-881951-2.
- Samson, Steven L. *MVS Performance Management*. New York, NY: McGraw-Hill (1997). ISBN 0-07-057700-5.
- Taylor, Ed. *Demystifying SNA*. Plano, TX: Wordware (1993). 1-55622-404-4.
- . *Demystifying TCP/IP*. Plano, TX: Wordware (1993). 1-55622-400-1.
- Terplan, Kornel, and Jill Huntington-Lee. *Distributed Systems and Network Management*. New York, NY: Van Nostrand Reinhold (1995). ISBN 0-442-01873-8.
- Young, John. *Exploring IBM's New Age Mainframes*. 5th ed. Gulf Breeze, FL: Maximum Press (1998). ISBN 978-1885068156.

11. Database Performance

Database performance focuses on tuning and optimizing the design, parameters, and physical construction of database objects, specifically tables and indexes, and the files in which their data is stored. The actual composition and structure of database objects must be monitored continually and changed accordingly if the database becomes inefficient. No amount of SQL tweaking or system tuning can optimize the performance of queries run against a poorly designed or disorganized database.

No amount of SQL tweaking or system tuning can optimize the performance of queries run against a poorly designed database.

Techniques for Optimizing Databases

The DBA must be cognizant of the features of the DBMS in order to apply the proper techniques for optimizing the performance of database structures. Most of the major DBMSs support the following techniques although perhaps by different names. Each of the following techniques can be used to tune database performance and will be discussed in subsequent sections.

- *Partitioning*—breaking a single database table into sections stored in multiple files
- *Raw partition versus file system*—choosing whether to store database data in an OS-controlled file or not
- *Indexing*—choosing the proper indexes and options to enable efficient queries
- *Denormalization*—varying from the logical design to achieve better query performance
- *Clustering*—enforcing the physical sequence of data on disk
- *Interleaving data*—combining data from multiple tables into a single, sequenced file
- *Free space*—leaving room for data growth
- *Compression*—algorithmically reducing storage requirements

- *File placement and allocation*—putting the right files in the right place
- *Page size*—using the proper page size for efficient data storage and I/O
- *Reorganization*—removing inefficiencies from the database by realigning and restructuring database objects

Partitioning

A database table is a logical manifestation of a set of data that physically resides on computerized storage. One of the decisions that the DBA must make for every table is how to store that data. Each DBMS provides different mechanisms that accomplish the same thing—mapping physical files to database tables. The DBA must decide from among the following mapping options for each table:

One decision that the DBA must make for every table is how to store its data.

- *Single table to a single file.* This is, by far, the most common choice. The data in the file is formatted such that the DBMS understands the table structure, and every row inserted into that table is stored in the same file. However, this setup is not necessarily the most efficient.
- *Single table to multiple files.* This option is used most often for very large tables or tables requiring data to be physically separated at the storage level. Mapping to multiple files is accomplished by using partitioned tablespaces or by implementing segmented disk devices.
- *Multiple tables to a single file.* This type of mapping is used for small tables, such as lookup tables and code tables, and can be more efficient from a disk utilization perspective.

Partitioning helps to accomplish parallelism. *Parallelism* is the process of using multiple tasks to access the database in parallel. A parallel request can be invoked to use multiple, simultaneous read engines for a single SQL statement. Parallelism is desirable because it can substantially reduce the elapsed time for database queries.

Partitioning helps to accomplish parallelism.

Multiple types of parallelism are based on the resources that can be invoked in parallel. For example, a single query can be broken down into

multiple requests, each using a different CPU core¹ in parallel. In addition, parallelism can be improved by spreading the work across multiple database instances. Each DBMS offers different levels of support for parallel database queries. To optimize database performance, the DBA should be cognizant of the support offered in each DBMS being managed and exploit the parallel query capabilities.

Raw Partition versus File System

For a UNIX-based DBMS environment,² the DBA must choose between a raw partition and using the UNIX file system to store the data in the database. A *raw partition* is the preferred type of physical device for database storage because writes are cached by the operating system when a file system is used. When writes are buffered by the operating system, the DBMS does not know whether the data has been physically copied to disk or not. When the DBMS cache manager attempts to write the data to disk, the operating system may delay the write until later because the data may still be in the file system cache. If a failure occurs, data in a database using the file system for storage may not be 100 percent recoverable. This is to be avoided.

If a raw partition is used instead, the data is written directly from the database cache to disk with no intermediate file system or operating system caching, as shown in [Figure 11.1](#). When the DBMS cache manager writes the data to disk, it will physically be written to disk with no intervention.




Figure 11.1. Using raw partitions to avoid file system caching

Additionally, when using a raw partition, the DBMS will ensure that enough space is available and write the allocation pages. When using a file system, the operating system will not preallocate space for database usage.

From a performance perspective, there is no advantage to having a secondary layer of caching at the file system or operating system level; the

DBMS cache is sufficient. Actually, the additional work required to cache the data a second time consumes resources, thereby negatively impacting the overall performance of database operations.

From a performance perspective, the DBMS cache is sufficient.

Indexing

Creating the correct indexes on tables in the database is perhaps the single greatest performance tuning technique that a DBA can perform. Indexes are used to enhance performance. They are particularly useful for

- Locating rows by value(s) in column(s)
- Making joins more efficient (when the index is defined on the join columns)
- Correlating data across tables
- Aggregating data
- Sorting data to satisfy a query

Indexes are used to enhance performance.

Without indexes, all access to data in the database would have to be performed by scanning all available rows. Scans are very inefficient for very large tables.

Designing and creating indexes for database tables actually crosses the line between database performance tuning and application performance tuning. Indexes are database objects created by the DBA with database DDL. However, an index is built to make SQL statements in application programs run faster. Indexing as a tuning effort is applied to the database to make applications more efficient when the data access patterns of the application vary from what was anticipated when the database was designed.

Before tuning the database by creating new indexes, be sure to understand the impact of adding an index. The DBA should have an understanding of the access patterns of the table on which the index will be built. Useful information includes the percentage of queries that access rather than update the table, the performance thresholds set within any service-level agreements for queries on the table, and the impact of adding a new index on running

database utilities such as loads, reorganizations, and recovery.

One of the big unanswered questions of database design is “How many indexes should be created for a single table?” There is no set answer to this question. DBAs will need to use their expertise to determine the proper number of indexes for each table such that database queries are optimized and the performance of database inserts, updates, and deletes does not degrade. Determining the proper number of indexes for each table requires in-depth analysis of the database and the applications that access the database.

The general goal of index analysis is to use less I/O to the database to satisfy the queries made against the table. Of course, an index can help some queries and hinder others. Therefore, the DBA must assess the impact of adding an index to all applications and not just tune single queries in a vacuum. This can be an arduous but rewarding task.

An index affects performance positively when fewer I/Os are used to return results to a query. Conversely, an index negatively impacts performance when data is updated and the indexes have to be changed as well. An effective indexing strategy seeks to provide the greatest reduction in I/O with an acceptable level of effort to keep the indexes updated.

Some applications have troublesome queries that require significant tuning to achieve satisfactory performance. Creating an index to support a single query is acceptable if that query is important enough in terms of ROI to the business (or if it is run by your boss or the CEO). If the query is run infrequently, consider creating the index before the process begins and dropping the index when the process is complete.

Whenever you create new indexes, be sure to thoroughly test the performance of the queries it supports. Additionally, be sure to test database modification statements to gauge the additional overhead of updating the new indexes. Review the CPU time, elapsed time, and I/O requirements to assure that the indexes help. Keep in mind that tuning is an iterative process, and it may take time and several index tweaks to determine the impact of a change. There are no hard-and-fast rules for index creation. Experiment with different index combinations and measure the results.

Be sure to thoroughly test the performance of the queries the index supports.

For additional information on the structure of indexes and the types of indexes commonly supported by DBMSs, refer to the section on designing indexes in [Chapter 4](#).

When to Avoid Indexing

There are a few scenarios where indexing may not be a good idea. When tables are very small, say, less than ten pages, consider avoiding indexes. Indexed access to a small table can be less efficient than simply scanning all of the rows because reading the index adds I/O requests. For example, SQL Server will always perform a table scan on tables of less than 64K (8 pages) in size.

Index I/O notwithstanding, even a small table can sometimes benefit from being indexed—for example, to enforce uniqueness or if most data access retrieves a single row using the primary key.

You may want to avoid indexing variable-length columns if the DBMS in question expands the variable column to the maximum length within the index. Such expansion can cause indexes to consume an inordinate amount of disk space and might be inefficient. However, if variable-length columns are used in SQL WHERE clauses, the cost of disk storage must be compared to the cost of scanning. Buying some extra disk storage is usually cheaper than wasting CPU resources to scan rows. Furthermore, the SQL query might contain alternate predicates that could be indexed instead of the variable-length columns. Some DBMSs support index compression, which can be used to reduce storage requirements for indexes. The pros and cons of compressing indexes are similar to those already discussed for table and tablespace compression.

It is also a good idea to avoid indexing a column with a very small number of values, often referred to as cardinality. When there are only a few values, an index is not very helpful and can become a bottleneck to maintain. GENDER is an example of a column with a low number of distinct values.

Additionally, avoid indexing any table that is always accessed using a scan; that is, the SQL issued against the table never supplies a WHERE clause.

Avoid indexing any table that is always accessed using a scan.

Index Overloading

Query performance can be enhanced in certain situations by overloading an index with additional columns. Indexes are typically based on the WHERE clauses of SQL SELECT statements. For example, consider the following SQL statement:

```
select emp_no, last_name, salary
from   employee
where  salary > 15000.00;
```

Creating an index on the salary column can enhance the performance of this query. However, the DBA can further enhance the performance of the query by overloading the index with the emp_no and last_name columns as well. With an overloaded index, the DBMS can satisfy the query by using *only* the index. The DBMS need not incur the additional I/O of accessing the table data, since every piece of data that is required by the query exists in the overloaded index.

DBAs should consider overloading indexes to encourage index-only access when multiple queries can benefit from the index or when individual queries are very important.

Denormalization

Another way to optimize the performance of database access is to denormalize the tables. Denormalization was covered in detail in [Chapter 4](#), so we will not discuss the topic in depth here. Suffice it to say, [denormalization](#), the opposite of [normalization](#), is the process of putting one fact in many places. This speeds data retrieval at the expense of data modification. Denormalizing tables can be a good decision when a completely normalized design does not perform optimally.

The only reason to denormalize a relational database is to enhance performance.

The only reason to ever denormalize a relational database design is to enhance performance. As discussed in [Chapter 4](#), you should consider the following options:

- *Prejoined tables*—when the cost of joining is prohibitive
- *Report table*—when specialized critical reports are too costly to generate

- *Mirror table*—when tables are required concurrently by two types of environments
- *Split tables*—when distinct groups use different parts of a table
- *Combined tables*—to consolidate one-to-one or one-to-many relationships into a single table
- *Speed table*—to support hierarchies like bills of materials or reporting structures
- *Physical denormalization*—to take advantage of specific DBMS characteristics

You might also consider

- Storing *redundant data* in tables to reduce the number of table joins required
- Storing *repeating groups* in a row to reduce I/O and possibly disk space
- Storing *derivable data* to eliminate calculations and costly algorithms

Clustering

A clustered table will store its rows physically on disk in order by a specified column or columns. Clustering usually is enforced by the DBMS with a clustering index. The clustering index forces table rows to be stored in ascending order by the indexed columns. The left-to-right order of the columns, as defined in the index, defines the collating sequence for the clustered index. There can be only one clustering sequence per table (because physically the data can be stored in only one sequence).

A clustered table will store its rows physically on disk in order by a specified column or columns.

[Figure 11.2](#) demonstrates the difference between clustered and unclustered data and indexes; the clustered index is on top, and the unclustered index is on the bottom. As you can see, the entries on the leaf pages of the top index are in sequential order—in other words, they are clustered. Clustering enhances the performance of queries that access data sequentially because fewer I/Os need to be issued to retrieve the same data.

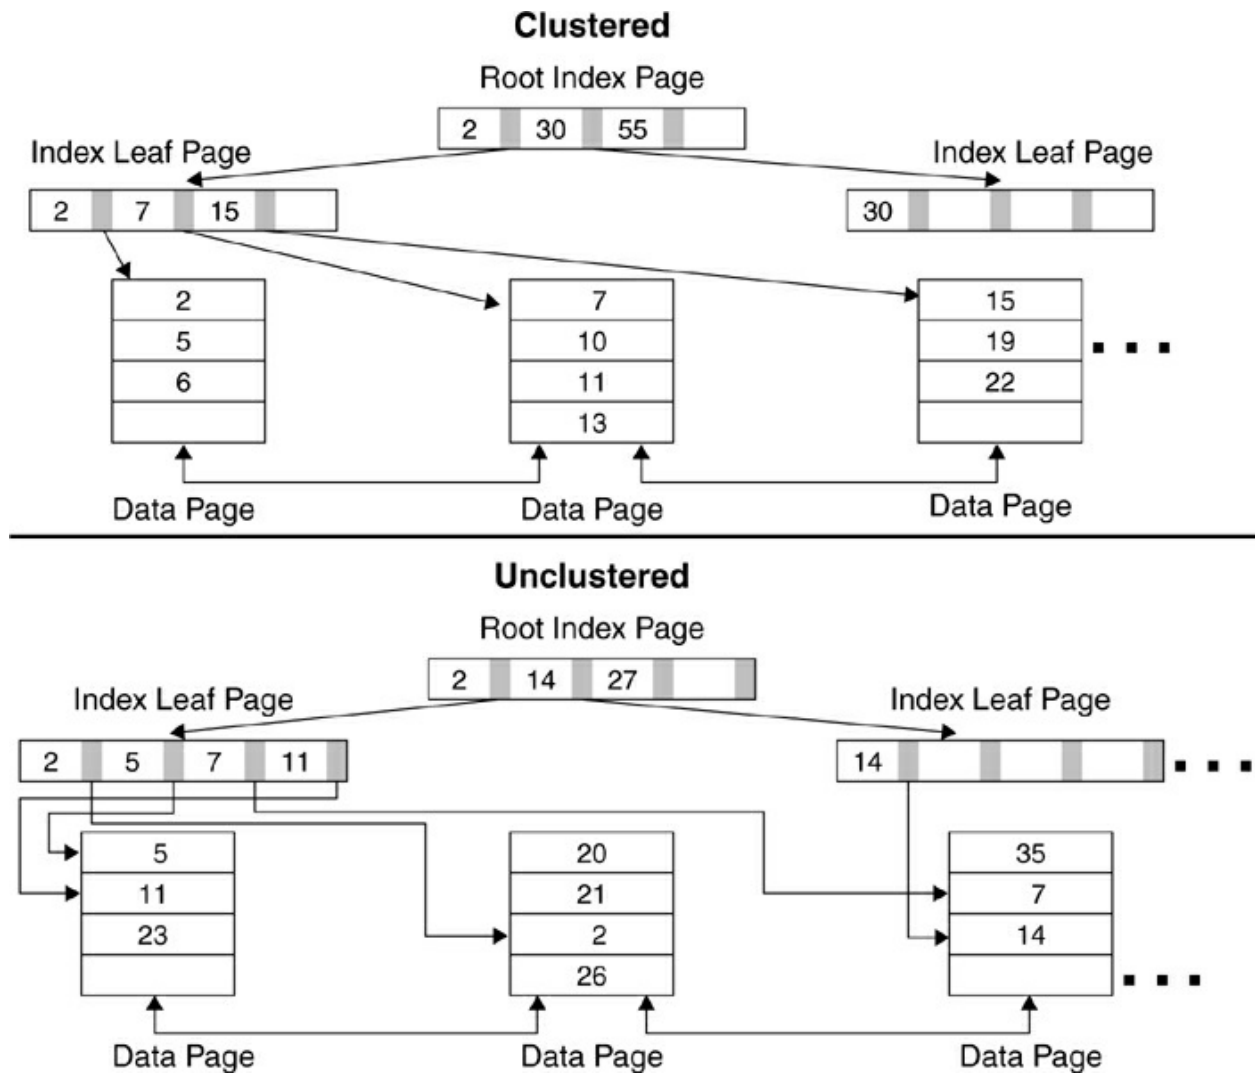


Figure 11.2. Clustered and unclustered indexes

Depending on the DBMS, the data may not always be physically maintained in exact clustering sequence. When a clustering sequence has been defined for a table, the DBMS will act in one of two ways to enforce clustering:

1. When new rows are inserted, the DBMS will physically maneuver data rows and pages to fit the new rows into the defined clustering sequence; or
2. When new rows are inserted, the DBMS will try to place the data into the defined clustering sequence, but if space is not available on the required page the data may be placed elsewhere.

The DBA must learn how the DBMS maintains clustering. If the DBMS

operates as in the second scenario, data may become unclustered over time and require reorganization. A detailed discussion of database reorganization appears later in this chapter. For now, though, back to our discussion of clustering.

Clustering tables that are accessed sequentially is good practice. In other words, clustered indexes are good for supporting range access, whereas unclustered indexes are better for supporting random access. Be sure to choose the clustering columns wisely. Use clustered indexes for the following situations:

- Join columns, to optimize SQL joins where multiple rows match for one or both tables participating in the join
- Foreign key columns, because they are frequently involved in joins and the DBMS accesses foreign key values during declarative referential integrity checking
- Predicates in a WHERE clause
- Range columns
- Columns that do not change often (reduces physically reclustered)
- Columns that are frequently grouped or sorted in SQL statements

Clustering tables that are accessed sequentially is good practice.

In general, the clustering sequence that aids the performance of the most commonly accessed predicates should be used for clustering. When a table has multiple candidates for clustering, weigh the cost of sorting against the performance gained by clustering for each candidate key. As a rule of thumb, though, if the DBMS supports clustering, it is usually a good practice to define a clustering index for each table that is created (unless the table is very small).

Clustering is generally not recommended for primary key columns because the primary key is, by definition, unique. However, if ranges of rows frequently are selected and ordered by primary key value, a clustering index may be beneficial.

Page Splitting

When the DBMS has to accommodate inserts, and no space exists, it must

create a new page within the database to store the new data. The process of creating new pages to store inserted data is called *page splitting*. A DBMS can perform two types of page splitting: *normal page splits* and *monotonic page splits*. Some DBMSs support both types of page splitting, while others support only one type. The DBA needs to know how the DBMS implements page splitting in order to optimize the database.

The process of creating new pages to store inserted data is called *page splitting*.

[Figure 11.3](#) depicts a normal page split. To accomplish this, the DBMS performs the following tasks in sequence:

1. Creates a new empty page between the full page and the next page
2. Takes half of the entries from the full page and moves them to the empty page
3. Adjusts any internal pointers to both pages and inserts the row accordingly

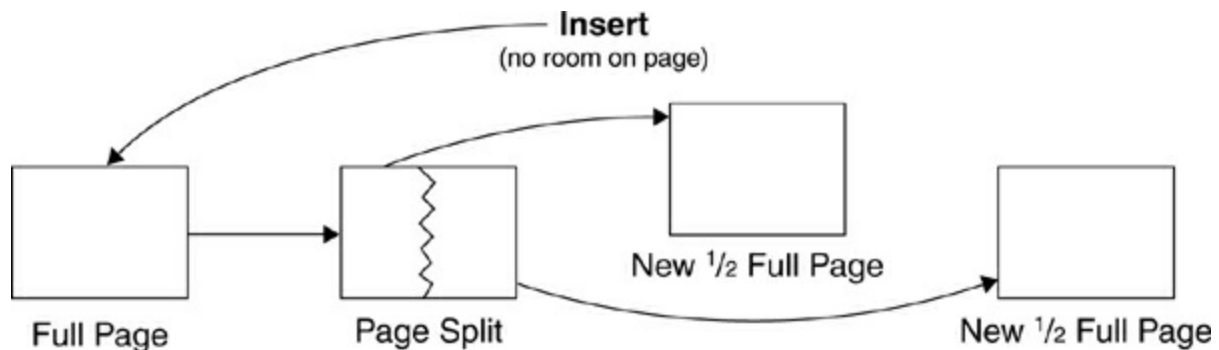


Figure 11.3. Normal page splitting

A monotonic page split is a much simpler process, requiring only two steps. The DBMS

1. Creates a new page between the full page and the next page
2. Inserts the new values into the fresh page

Monotonic page splits are useful when rows are being inserted in strictly ascending sequence. Typically, a DBMS that supports monotonic page splits will invoke it when a new row is added to the end of a page and the last addition was also to the end of the page.

When ascending rows are inserted and normal page splitting is used, a lot

of space can be wasted because the DBMS will be creating half-full pages that never fill up. If the wrong type of page split is performed during database processing, wasted space will ensue, requiring the database object to be reorganized for performance.

Interleaving Data

When data from two tables is frequently joined, it can make sense to physically interleave the data into the same physical storage structure. This can be viewed as a specialized form of clustering (and, in fact, Oracle uses the term *cluster* to define interleaved data). Interleaving data is covered in [Chapter 4](#) and is mentioned here as a performance tuning technique to consider.

Interleaving can be viewed as a specialized form of clustering.

Free Space

Free space, sometimes called *fill factor*, can be used to leave a portion of a tablespace or index empty and available to store newly added data. The specification of free space in a tablespace or index can reduce the frequency of reorganization, reduce contention, and increase the efficiency of insertion. Each DBMS provides a method of specifying free space for a database object in the CREATE and ALTER statements. A typical parameter is PCTFREE, where the DBA specifies the percentage of each data page that should remain available for future inserts. Another possible parameter is FREEPAGE, where the DBA indicates the specified number of pages after which a completely empty page is available.

Ensuring a proper amount of free space for each database object provides the following benefits:

- Inserts are faster when free space is available.
- As new rows are inserted, they can be properly clustered.
- Variable-length rows and altered rows have room to expand, potentially reducing the number of relocated rows.
- Having fewer rows on a page results in better concurrency because less data is unavailable to other users when a page is locked.

Ensure a proper amount of free space for each database object.

However, free space also has several disadvantages:

- Disk storage requirements are greater.
- Scans take longer.
- Fewer rows on a page can necessitate more I/O operations to access the requested information.
- Because the number of rows per page decreases, the efficiency of data caching can decrease because fewer rows are retrieved per I/O.

The DBA should monitor free space and ensure that the appropriate amount is defined for each database object. The correct amount of free space must be based on

- Frequency of inserts and modifications
- Amount of sequential versus random access
- Impact of accessing unclustered data
- Type of processing
- Likelihood of row chaining, row migration, and page splits

Don't define a static table with free space—it will not need room in which to expand.

Compression

Compression can be used to shrink the size of a database. When data is compressed, the database requires less disk storage. Some DBMSs provide internal DDL options to compress database files; third-party software is available for those that do not provide such features.

When compression is specified, data is algorithmically compressed upon insertion into the database and decompressed when it is read. Reading and writing compressed data consumes more CPU resources than reading and writing uncompressed data: The DBMS must execute code to compress and decompress the data as users insert, update, and read the data.

So why compress data? Consider an uncompressed table with a row size of 800 bytes. Five of this table's rows would fit in a 4K data page (or block). Now what happens if the data is compressed? Assume that the compression routine achieves 30 percent compression on average (a very conservative estimate). In that case, the 800-byte row will consume only 560 bytes ($800 \times 0.30 = 560$). After the data is compressed, seven rows will fit on a 4K page.

Because I/O occurs at the page level, a single I/O will retrieve more data, which will optimize the performance of sequential data scans and increase the likelihood of data residing in the cache because more rows fit on a physical page.

Of course, compression always requires a trade-off that the DBA must analyze. On the positive side, we have disk savings and the potential for reducing I/O cost. On the negative side, we have the additional CPU cost required to compress and decompress the data.

Compression always requires a trade-off, but in I/O-heavy systems it is usually worth it.

However, compression is not an option for every database index or table. For smaller amounts of data, it is possible that a compressed file will be larger than an uncompressed file. This is so because some DBMSs and compression algorithms require an internal dictionary to manage the compression. The dictionary contains statistics about the composition of the data that is being compressed. For a trivial amount of data, the size of the dictionary may be greater than the amount of storage saved by compression.

File Placement and Allocation

The location of the files containing the data for the database can have an impact on performance. A database is very I/O intensive, and the DBA must make every effort to minimize the cost of physical disk reading and writing.

The DBA must make every effort to minimize the cost of physical disk reading and writing.

This discipline entails

- Understanding the access patterns associated with each piece of data in the system
- Placing the data on physical disk devices in such a way as to optimize performance

The first consideration for file placement on disk is to separate the indexes from the data, if possible. Database queries are frequently required to access data from both the table and an index on that table. If both of these files reside on the same disk device, performance degradation is likely. To retrieve

data from disk, an arm moves over the surface of the disk to read physical blocks of data on the disk. If a single operation is accessing data from files on the same disk device, latency will occur; reads from one file will have to wait until reads from the other file are processed. Of course, if the DBMS combines the index with the data in the same file, this technique cannot be used.

Another rule for file placement is to analyze the access patterns of your applications and separate the files for tables that are frequently accessed together. DBAs should do this for the same reason they should separate index files from table files.

A final consideration for placing files on separate disk devices occurs when a single table is stored in multiple files (partitioning). It is wise in this case to place each file on a separate disk device to encourage and optimize parallel database operations. If the DBMS can break apart a query to run it in parallel, placing multiple files for partitioned tables on separate disk devices will minimize disk latency.

With modern disk systems such as RAID devices, precise file placement is often difficult, if not impossible, to achieve. RAID storage consists of arrays of disk drives, and the files stored on the device are split across multiple physical disk drives. Be sure to read the documentation for your specific DBMS and follow the guidance for the placement of critical system files (such as log and configuration files). And keep in mind that a hard disk drive with two partitions, one holding log files and the other partition holding data files, is not the same as two disks, each dedicated to either log or data files.

For more details on storage and RAID please refer to [Chapter 18, “Data and Storage Management.”](#)

Database Log Placement

Placing the transaction log on a separate disk device from the actual data allows the DBA to back up the transaction log independently from the database. It also minimizes dual writes to the same disk. Writing data to two files on the same disk drive at the same time will degrade performance even more than reading data from two files on the same disk drive at the same time. Remember, too, that every database modification (write) is recorded on the database transaction log.

Distributed Data Placement

The goal of data placement is to optimize access by reducing contention on physical devices. Within a client/server environment, this goal can be expanded to encompass the optimization of application performance by reducing network transmission costs.

Data placement optimizes access by reducing contention on physical devices.

Data should reside at the database server where it is most likely, or most often, to be accessed. For example, Chicago data should reside at the Chicago database server, Los Angeles–specific data should reside at the Los Angeles database server, and so on. If the decision is not so clear-cut (e.g., San Francisco data if there is no database server in San Francisco), place the data on the database server that is geographically closest to where it will be most frequently accessed (in the case of San Francisco, L.A., not Chicago).

Be sure to take fragmentation, replication, and snapshot tables into account when deciding upon the placement of data in your distributed network.

Disk Allocation

The DBMS may require disk devices to be allocated for database usage. If this is the case, the DBMS will provide commands to initialize physical disk devices. The disk initialization command will associate a logical name for a physical disk partition or OS file. After the disk has been initialized, it is stored in the system catalog and can be used for storing table data.

Before initializing a disk, verify that sufficient space is available on the physical disk device. Likewise, make sure that the device is not already initialized.

Use meaningful device names to facilitate more efficient usage and management of disk devices. For example, it is difficult to misinterpret the usage of a device named DUMP_DEV1 or TEST_DEV7. However, names such as XYZ or A193 are not particularly useful. Additionally, maintain documentation on initialized devices by saving script files containing the actual initialization commands and diagrams indicating the space allocated by device.

Use meaningful device names to facilitate more efficient usage.

Page Size (Block Size)

Most DBMSs provide the ability to specify a page, or block, size. The *page size* is used to store table rows (or, more accurately, records that contain the row contents plus any overhead) on disk. For example, consider a table requiring rows that are 125 bytes in length with 6 additional bytes of overhead. This makes each record 131 bytes long. To store 25 records on a page, the page size would have to be at least 3,275 bytes. However, each DBMS requires some amount of page overhead as well, so the practical size will be larger. If page overhead is 20 bytes, the page size would be 3,295—that is, $3,275 + 20$ bytes of overhead.

Choosing the proper page size is an important DBA task.

This discussion, however, is simplistic. In general practice, most tablespaces will require some amount of free space to accommodate new data. Therefore, some percentage of free space will need to be factored into the equation.

To complicate matters, many DBMSs limit the page sizes that can be chosen. For example, DB2 for z/OS limits page size to 4K, 8K, 16K, or 32K; Microsoft SQL Server supports only an 8K page size. In the case of DB2, the DBA will need to calculate the best page size based on row size, the number of rows per page, and free space requirements.

Consider this question: “In DB2 for z/OS, what page size should be chosen if 0 percent free space is required and the record size is 2,500 bytes?”

The simplistic answer is 4K, but it might not be the best answer. A 4K page would hold one 2,500-byte record per page, but an 8K page would hold three 2,500-byte records. The 8K page could provide for more efficient sequential I/O processing, because reading 8K of data would return three rows, whereas reading 8K of data using two 4K pages would return only two rows.

Choosing the proper page size is an important DBA task for optimizing database I/O performance.

Database Reorganization

Relational technology and SQL make data modification easy. Just issue an INSERT, UPDATE, or DELETE statement with the appropriate WHERE

clause and the DBMS takes care of the actual data navigation and modification. In order to provide this level of abstraction, the DBMS handles the physical placement and movement of data on disk. Theoretically, this makes everyone happy. The programmer's interface is simplified, and the RDBMS takes care of the hard part—manipulating the actual placement of data. However, things are not quite that simple. The manner in which the DBMS physically manages data can cause subsequent performance problems.

Every DBA has encountered the situation where a query or application that used to perform well slows down after it has been in production for a while. These slowdowns have many potential causes—perhaps the number of transactions issued has increased, or the volume of data has expanded. However, the performance problem might be due to database disorganization. Database disorganization occurs when a database's logical and physical storage allocations contain many scattered areas of storage that are too small, not physically contiguous, or too disorganized to be used productively. Let's review the primary culprits:

Query or application slowdowns have many potential causes.

- The first possibility is *unclustered data*. If the DBMS does not strictly enforce clustering, a clustered table or index can become unclustered as data is added and changed. If the data becomes significantly unclustered, the DBMS cannot rely on the clustering sequence. Because the data is no longer clustered, queries that were optimized to access data cannot take advantage of the clustering sequence. In this case, the performance of queries run against the unclustered table will suffer.
- *Fragmentation* is a condition in which there are many scattered areas of storage in a database that are too small to be used productively. It results in wasted space, which can hinder performance because additional I/Os are required to retrieve the same data.
- *Row chaining* or *row migration* occurs when updated data does not fit in the space it currently occupies, and the DBMS must find space for the row. With row chaining, the DBMS moves a part of the new, larger row to a location within the tablespace where free space exists. With row migrations, the full row is placed elsewhere in the tablespace. In each case, a pointer is used to locate either the rest of the row or the full row. Both row chaining and row migration will result in the issuance of

multiple I/Os to read a single row. Performance will suffer because multiple I/Os are more expensive than a single I/O.

- *Page splits* can cause disorganized databases, too. If the DBMS performs monotonic page splits when it should perform normal page splits, or vice versa, space may be wasted. When space is wasted, fewer rows exist on each page, causing the DBMS to issue more I/O requests to retrieve data. Therefore, once again, performance suffers.
- *File extents* can negatively impact performance. An *extent* is an additional file that is tied to the original file and can be used only in conjunction with the original file. When the file used by a tablespace runs out of space, an extent is added for the file to expand. However, file extents are not stored contiguously with the original file. As additional extents are added, data requests will need to track the data from extent to extent, and the additional code this requires is unneeded overhead. Resetting the database space requirements and reorganizing can clean up file extents.

Let's take a look at a disorganized tablespace by comparing [Figures 11.4](#) and [11.5](#). Assume that a tablespace consists of three tables across multiple blocks, such as the tablespace and tables depicted in [Figure 11.4](#). Each box represents a data page.

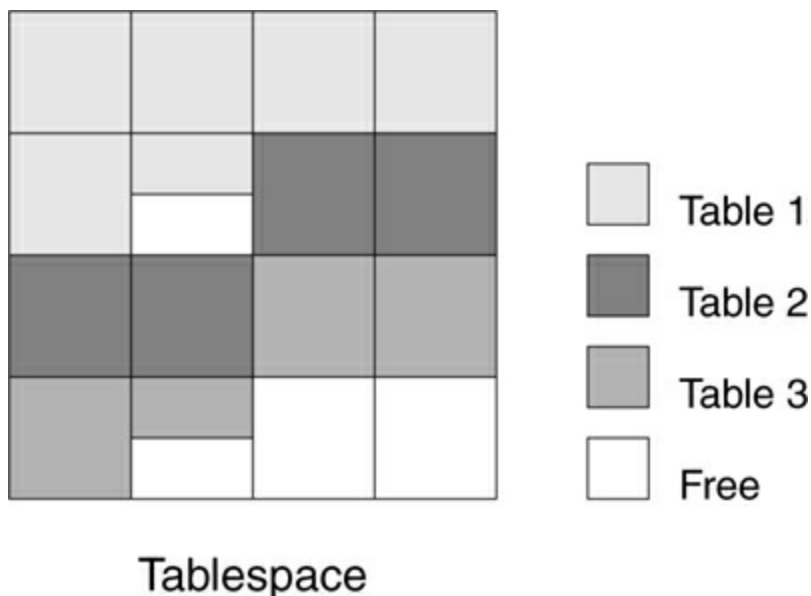


Figure 11.4. Organized tablespace

Now, let's make a couple of changes to the data in these tables. First, we'll

add six rows to the second table. However, no free space exists into which these new rows can be stored. How can the rows be added? The DBMS requires an additional extent to be taken into which the new rows can be placed. This results in fragmentation: The new rows have been placed in a noncontiguous space. For the second change, let's update a row in the first table to change a variable-length column; for example, let's change the value in a LASTNAME column from WATSON to BEAUCHAMP. Issuing this update results in an expanded row size because the value for LASTNAME is longer in the new row: "BEAUCHAMP" contains nine characters whereas "WATSON" consists of only six. This action results in row chaining. The resultant tablespace shown in [Figure 11.5](#) depicts both the fragmentation and the row chaining.




Figure 11.5. Disorganized tablespace

Depending on the DBMS, there may be additional causes of disorganization. For example, if multiple tables are defined within a tablespace, and one of the tables is dropped, the tablespace may need to be reorganized to reclaim the space.

To correct disorganized database structures, the DBA can run a database or tablespace reorganization utility, or REORG, to force the DBMS to restructure the database object, thus removing problems such as unclustered data, fragmentation, and row chaining. The primary benefit of reorganization is the resulting speed and efficiency of database functions because the data is

organized in a more optimal fashion on disk. In short, reorganization maximizes availability and reliability for databases.

Running a reorganization causes the DBMS to restructure the database object, maximizing the availability, speed, and efficiency of database functions.

Both tablespaces and indexes can be reorganized. How the DBA runs a REORG utility depends on the DBMS. Some DBMS products ship with a built-in reorganization utility. Others require the customer to purchase the utility. Still others claim that the customer will not need the utility at all when using their DBMS. I have found the last claim to be untrue. Every DBMS incurs some degree of disorganization as data is added and modified.

Of course, DBAs can manually reorganize a database by completely rebuilding it. However, accomplishing such a reorganization requires a complex series of steps. [Figure 11.6](#) depicts the steps entailed by a manual reorganization.

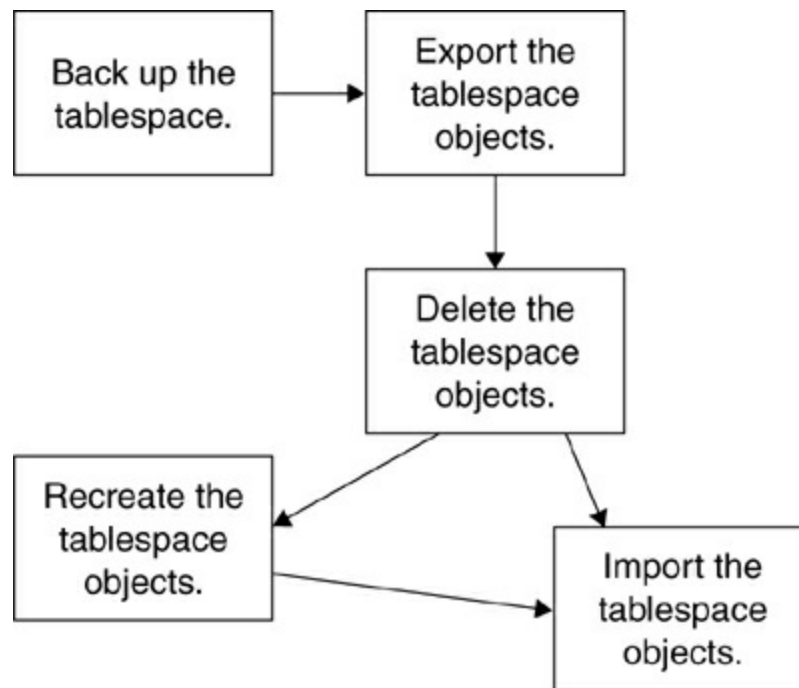


Figure 11.6. Typical steps for a manual reorganization

If a utility is available for reorganizing, from either the DBMS vendor or a third-party vendor, the process is greatly simplified. Sometimes the utility is as simple as issuing a command such as

REORG TABLESPACE TSNAME

A traditional reorganization requires the database to be down. The high cost of downtime creates pressure both to perform and to delay preventive maintenance—a no-win situation familiar to most DBAs. Some REORG utilities are available that perform the reorganization while the database is online. Such a reorganization is accomplished by making a copy of the data. The online REORG utility reorganizes the copy while the original data remains online. When the copied data has been reorganized, an online REORG uses the database log to “catch up” by applying to the copy any data changes that occurred during the process. When the copy has caught up to the original, the online REORG switches the production tablespace from the original to the copy. Performing an online reorganization requires additional disk storage and a slow transaction window. If a large number of transactions occur during the online reorganization, REORG may have a hard time catching up.

Determining When to Reorganize

System catalog statistics can help to determine when to reorganize a database object. Each DBMS provides a method of reading through the contents of the database and recording statistical information about each database object. Depending on the DBMS, this statistical information is stored either in the system catalog or in special pages within the database object itself.

One statistic that can help a DBA determine when to reorganize is *cluster ratio*. Cluster ratio is the percentage of rows in a table that are actually stored in clustering sequence. The closer the cluster ratio is to 100 percent, the more closely the actual ordering of the rows on the data pages matches the clustering sequence. A low cluster ratio indicates bad clustering, and a reorganization may be required. A low cluster ratio, however, may not be a performance hindrance if the majority of queries access data randomly instead of sequentially.

Cluster ratio is the percentage of rows in a table that are actually stored in clustering sequence.

Tracking down the other causes of disorganization can sometimes be difficult. Some DBMSs gather statistics on fragmentation, row chaining, row migration, space dedicated to dropped objects, and page splits; others do not.

Oracle provides a plethora of statistics in dynamic performance tables that can be queried. Refer to the sidebar “[Oracle Dynamic Performance Tables](#)” for more details.³

Tablespaces are not the only database objects that can be reorganized. Indexes, too, can benefit from reorganization. As table data is added and modified, the index too must be changed. Such changes can cause the index to become disorganized.

Oracle Dynamic Performance Tables

Oracle stores vital performance statistics about the database system in a series of dynamic performance tables. These tables are sometimes referred to as the “V\$ tables” because the table names are prefixed with the characters V\$.

The V\$ tables are used by the built-in Oracle performance monitoring facilities and can be queried by the DBA for insight into the well-being and performance of an Oracle instance. Examples of some of the statistics that can be found in the V\$ tables include

- Free space available
- Chained rows
- Rollback segment contention and activity
- Memory usage
- Lock and latch activity
- Buffer waits
- Disk activity

Of course, there is quite a lot of additional performance information to be found in these tables. Oracle DBAs should investigate the V\$ tables and query them regularly to analyze the performance of the Oracle system, its databases, and applications.

A vital index statistic to monitor is the number of levels. Recall from [Chapter 4](#) that most relational indexes are b-tree structures. As data is added to the index, the number of levels of the b-tree will grow. When more levels exist in the b-tree, more I/O requests are required to move from the top of the index structure to the actual data that must be accessed. Reorganizing an

index can cause the index to be better structured and require fewer levels.

Another index statistic to analyze to determine whether reorganization is required is the distance between the index leaf pages, or *leaf distance*. Leaf distance is an estimate of the average number of pages between successive leaf pages in the index. Gaps between leaf pages can develop as data is deleted from an index or as a result of page splitting. Of course, the best value for leaf distance is zero, but achieving a leaf distance of zero in practice is not realistic. In general, the lower this value, the better. Review the value over time to determine a high-water mark for leaf distance that will indicate when indexes should be reorganized.

Automation

If possible, the DBA should look into using the database utilities or third-party tools to automate reorganizations. The automation tool can be used to query the database statistics and trigger reorganization only for those database objects that have surpassed the high-water mark for a combination of statistics. For example, the DBA may want to automatically reorganize all tablespaces where the cluster ratio is below 85 percent and all indexes where the leaf distance has a value greater than 100.

Look into using the database utilities or third-party tools to automate reorganizations.

Reorganizations can be costly in terms of downtime and computing resources. Furthermore, determining when reorganization will actually create performance gains can be difficult. However, the performance gains that can be accrued with reorganization are tremendous when data is unclustered, fragmented, or disorganized in any way. Wise DBAs will plan for and schedule reorganizations to resolve disorganization problems in their database systems.

Summary

Applications and data are constantly changing. Users require instant response time and 24/7 availability. The database structures that support these applications must be properly maintained to ensure optimal application performance. Proper database design, appropriate choice of clustering, and reorganizing databases based on statistics help to deliver efficient databases.

Furthermore, DBAs can ensure database performance by automating these processes to reduce the risk and errors associated with manual database maintenance.

Review

1. What is the best performance tuning technique a DBA can use to improve database performance?
2. What is the only reason to denormalize a physical data model?
3. What is the benefit of clustering data?
4. How many indexes should be defined for a single table?
5. Discuss the pros and cons of providing database objects with free space.
6. What are the causes of database and index disorganization?
7. How can file extents degrade database performance?
8. What performance advantages can be gained by partitioning a table?
9. Explain what the term *leaf distance* means and why it is important for index performance.
10. What is the benefit of allocating tablespaces and indexes on separate disk devices?

Suggested Reading

Colledge, Rod. *SQL Server 2008 Administration in Action*. Greenwich, CT: Manning Publications (2010). ISBN 978-1-933988-72-6.

Dunham, Jeff. *Database Performance Tuning Handbook*. New York, NY: McGraw-Hill (1998). ISBN 0-07-018244-2.

Informix Software Staff. *Evolution of the High Performance Database*. Upper Saddle River, NJ: Informix Press/Prentice Hall (1997). ISBN 0-13-594730-8.

Kyte, Thomas. *Effective Oracle by Design*. New York, NY: McGraw-Hill (2003). ISBN 0-07-223065-7.

Lahdenmaki, Tapio, and Michael Leach. *Relational Database Index Design and the Optimizers*. Hoboken, NJ: John Wiley & Sons (2005). ISBN 0-471-71999-4.

- Lewis, Jonathan. *Oracle Core: Essential Internals for DBAs and Developers*. New York, NY: Apress (2011). ISBN 978-1-4302-3954-3.
- Mullins, Craig S. *DB2 Developer's Guide*. 6th ed. Boston, MA: IBM Press (2012). ISBN 978-0-13-283642-5.
- Purba, Sanjiv. *High-Performance Web Databases*. Boca Raton, FL: Auerbach (2001). ISBN 0-8493-0882-8.
- Schwartz, Baron, et al. *High Performance MySQL*. 2nd ed. Sebastopol, CA: O'Reilly (2008). ISBN 978-0-596-10171-8.

12. Application Performance

Now that we have discussed performance from the system and database perspective, it is time to turn our attention to performance from the application code perspective. This chapter focuses on tuning and optimizing the application code and SQL statements, as well as ensuring the application interacts with the DBMS appropriately and efficiently. Poorly constructed and formulated application code accounts for the majority of relational database performance problems—as many as 80 percent of performance problems can be traced back to the application.

Poorly constructed and formulated application code accounts for most relational database performance problems.

Designing Applications for Relational Access

Application design issues were discussed in [Chapter 5](#) and should be revisited when application performance suffers. Perhaps the application was improperly designed and portions of the application, or perhaps the entire application, will need to be rewritten. The application must be designed for performance at the outset—because changing the design of the application later is either impossible, impractical, or too costly to tackle.

Design issues to examine when application performance suffers include the following:

- *Type of SQL.* Is the correct type of SQL (planned or unplanned, dynamic or static, embedded or stand-alone) being used for this particular application?
- *Programming language.* Is the programming language capable of achieving the required performance, and is the language environment optimized for database access?
- *Transaction design and processing.* Are the transactions within the program properly designed to assure ACID properties, and does the program use the transaction processor of choice appropriately and efficiently?
- *Locking strategy.* Does the application hold the wrong type of locks, or

- does it hold the correct type of locks for too long?
- *COMMIT strategy*. Does each application program issue SQL COMMIT statements to minimize the impact of locking?
 - *Batch processing*. Are batch programs designed appropriately to take advantage of the sequential processing features of the DBMS?
 - *Online processing*. Are online applications designed to return useful information and to minimize the amount of information returned to the user's screen for a single invocation of the program?

Relational Optimization

The DBA must become intimately familiar with the optimization techniques used by each DBMS in the organization. Of course, application developers must code efficient SQL and understand how to optimize SQL, but in the end it is the DBA who is responsible for the performance of database applications. As such, the DBA must be skilled in SQL coding and tuning SQL for performance.

The *optimizer* is the heart of a relational database management system. It is an inference engine responsible for determining the best possible database navigation strategy for any given SQL request.

The optimizer is an inference engine for determining the database navigation strategy.

The application developer specifies *what* data is needed by coding the SQL statements, the DBMS supplies information about *where* the data is located, and the relational optimizer decides *how* to efficiently navigate the database. The end user needs no knowledge of where and how the actual data is stored. The optimizer knows this information.

To optimize SQL, the relational optimizer must analyze each SQL statement by parsing it to determine the tables and columns that must be accessed. The optimizer will also access statistics stored by the RDBMS in either the system catalog or the database objects themselves. The statistics are used to determine the best method of accomplishing the tasks that need to be performed to satisfy the SQL request. This process is called *relational optimization* and is shown in [Figure 12.1](#).

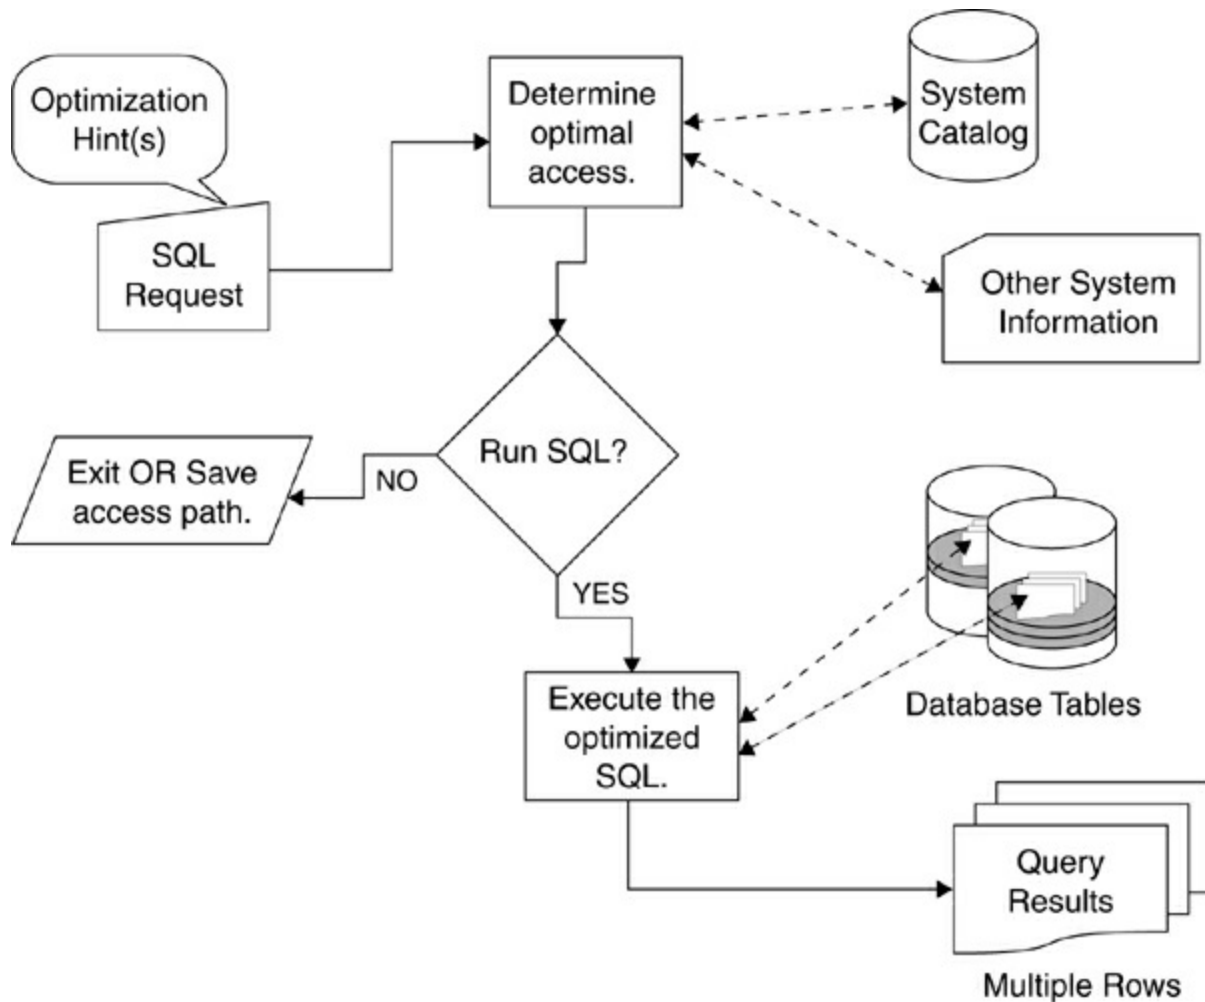


Figure 12.1. Relational optimization

Relational optimization is very powerful because it allows queries to adapt to a changing database environment. The optimizer can react to changes by formulating new access paths without requiring application coding changes to be implemented. The application can therefore be flexible as tables expand or contract in size, as indexes are added or removed, and as the database becomes disorganized or reorganized.

Regardless of how the data is physically stored and manipulated, SQL can be used to access data, and the DBMS will take the current state of the database into account to optimize that data access. This separation of access criteria from physical storage characteristics is called *physical data independence*.

The separation of access criteria from physical storage

characteristics is called physical data independence.

Every RDBMS has an embedded relational optimizer that renders SQL statements into executable access paths. Furthermore, each vendor's relational optimizer works a little differently, with different steps and using different information. Nevertheless, the core of the process is the same from DBMS to DBMS. The optimizer parses the SQL statement and performs various phases of optimization, typically involving verification of syntactic and semantic correctness, followed by query analysis and formulation of the access paths to satisfy the query.

The relational optimizer can deploy many types of strategies available to the DBMS for optimizing SQL statements. The internal operations and instructions that are used by each DBMS's optimizer are closely guarded secrets. Modern relational optimizers are cost based, meaning that the optimizer will attempt to formulate an access path for each query that reduces overall cost. To function in this manner, the optimizer must evaluate and analyze multiple factors, including estimated CPU and I/O costs, database statistics, and the actual SQL statement.

CPU and I/O Costs

A relational optimizer uses formulas and models to estimate the machine cost to run each potential access path for the query being optimized. Based on CPU information, the optimizer can arrive at a rough estimate of the CPU time required to run the query using each optimized access path it analyzes.

Furthermore, a relational optimizer must estimate the cost of the actual writing and retrieval of the data. The optimizer estimates the cost of I/O to the query by using a series of formulas based on the database statistics, the data cache efficiency, and the cost of I/O to intermediate work files. These formulas result in a *filter factor*, which determines the relative I/O cost of the query.

Database Statistics

A relational optimizer is of little use without accurate statistics about the data stored in the database. A relational DBMS provides a utility program or command to gather statistics about database objects and to store them for use by the optimizer (or by the DBA for performance monitoring). For example, to collect statistics in DB2, the DBA must execute the RUNSTATS utility; to

collect statistics in SQL Server the UPDATE STATISTICS command is issued.

A relational optimizer is of little use without accurate statistics.

The DBA should collect modified statistics whenever a significant volume of data has been added or modified. Failure to do so will result in the optimizer basing its cost estimates on inaccurate statistics. This may be detrimental to query performance.

Database statistics provide the optimizer with information about the state of the tablespaces, tables, columns, and indexes. The DBMS collects statistical information such as

- Number of rows in the tablespace, table, or index
- Number of unique values stored in the column
- Most frequently occurring values for columns
- Index key density (see the “[Density](#)” sidebar)

Density

Density is the average percentage of duplicate values stored in the index key column(s) and is recorded as a percentage. For example, a density of 50 percent means that for a given index key value, we would expect to return 50 percent of the rows. The following equation determines the average number of rows expected to be returned when accessing a table by the index:

$$\text{Average \# rows} = \text{Total \# rows} \times \text{Density}$$

For example, when accessing a table with 1,000 rows by an index with a density of 50 percent:

$$\text{Average \# rows} = 1,000 \times 0.50 = 500$$

When accessing a table with 20,000 rows by an index with a density of 15 percent:

$$\text{Average \# rows} = 20,000 \times 0.15 = 3,000$$

This information is useful to the optimizer because it assists in determining the size of the results set and thus whether an index is helpful for specific access paths.

- Details on the ratio of clustering for clustered tables
- Correlation of columns to other columns
- Structural state of the index or tablespace
- Amount of storage used by the database object

When issuing RUNSTATS or the UPDATE STATISTICS command, the DBA specifies which statistics to gather. Of course, the exact statistics collected vary from DBMS to DBMS; additional or fewer statistics may be available within your database system. The key, though, is to keep the statistics as accurate as possible to ensure efficient and useful relational optimization.

When developing an application against test databases, the statistics for the test data will not accurately reflect the statistics for the production database. Whenever possible, the DBA should work with the application development team to create a script to populate production statistics into the test system. Depending on the DBMS, this may be accomplished with SQL statements or a data testing tool. Without production statistics, the DBMS will likely choose different access paths in the test environment from the ones it would choose in production—potentially causing performance problems when the application goes into production status.

Create a script to populate production statistics into the test system.

Query Analysis

The *query analysis* scans the SQL statement to determine its overall complexity. The formulation of the SQL statement is a significant factor in determining the access paths chosen by the optimizer. The complexity of the query, the number and type of predicates, the presence of functions, and the presence of ordering clauses enter into the estimated cost that is calculated by the optimizer. The more complex the SQL statement, the more work the query analysis must do to make sense of the SQL statement. During query analysis, the optimizer analyzes aspects of the SQL statement and the

database system, such as

The more complex the SQL statement, the more work the query analysis must do.

- Which tables in which database are required
- Whether any views are required to be broken down into underlying tables
- Whether table joins or subselects are required
- Whether UNION, EXCEPT, or INTERSECT are required
- Which indexes, if any, can be used
- How many predicates (WHERE clauses) must be satisfied
- Which functions must be executed
- Whether the SQL uses OR or AND
- How the DBMS processes each component of the SQL statement
- How much memory has been assigned to the data cache(s) used by the tables in the SQL statement
- How much memory is available for sorting if the query requires a sort

In other words, the query analysis breaks down the SQL statement into discrete tasks that must be performed to return the query results.

A large part of query analysis is index selection. After the optimizer determines the indexes available to be used for each predicate, it will decide whether to use a single index, multiple indexes, or no index at all.

Joins

When multiple tables are accessed, the optimizer figures out how to combine the tables in the most efficient manner. Combining information from multiple tables is known as *joining*. When determining the access path for a join, the optimizer must determine the order in which the tables will be joined, compute the overall cost estimate of each access path, and choose a join method for the particular query. The DBMS can use several different methods for joining tables.

Combining information from multiple tables is known as joining.

Regardless of the join method, the DBMS must make several decisions and perform certain operations. The first decision is to choose the table to process first—this table is referred to as the *outer table*. Next, a series of operations are performed on the outer table to prepare it for joining. Rows from that table are then combined with rows from the second table, called the *inner table*, as shown in [Figure 12.2](#). A series of operations are also performed on the inner table before the join occurs, as the join occurs, or both.

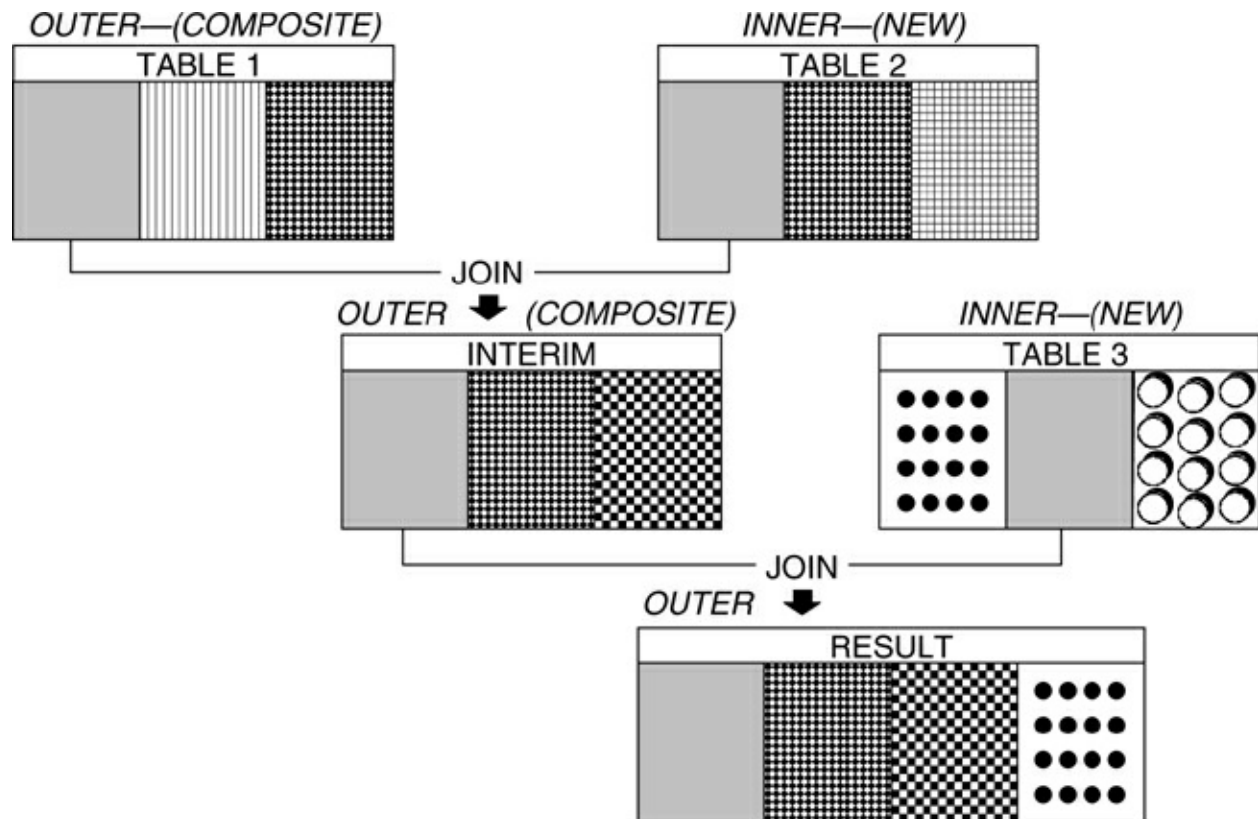


Figure 12.2. Joining tables

Although all joins are similar in functionality, each join method works differently behind the scenes. Let's investigate two common join methods: the *nested-loop join* and the *merge-scan join*. There are other join methods such as DB2's *hybrid join* and SQL Server's *hash join*, but we will discuss only the two most common methods here.

The *nested-loop join* works by comparing qualifying rows of the outer table to the inner table. A qualifying row is identified in the outer table, and then the inner table is scanned for a match. A *qualifying row* is one in which the predicates for columns in the table match. When the inner table scan is complete, another qualifying row in the outer table is identified. The inner

table is scanned for a match again, and so on. The repeated scanning of the inner table is usually accomplished with an index to avoid undue I/O costs. The smaller the size of the inner table, the better a nested-loop join performs, because fewer rows need to be scanned for each qualifying row of the outer table.

A second type of join method is the *merge-scan join*. In a merge-scan join, the tables to be joined are ordered by the keys. This ordering can be accomplished by a sort or by access via an index. After ensuring that both the outer and inner tables are properly sequenced, each table is read sequentially, and the join columns are matched. During a merge-scan join, no row from either table is read more than once. Merge-scan joins are useful when an appropriate index is not available on one (or both) of the tables.

Depending on the DBMS, other join methods may be available. For example, star joins to support data warehousing and analytical queries are typical in many DBMSs. Additional information on star schema and star join can be found in [Chapter 20, “Data Warehouse Administration.”](#)

Join Order

The optimizer reviews each join in a query and analyzes the appropriate statistics to determine the optimal order in which the tables should be accessed to accomplish the join. To find the optimal join access path, the optimizer uses built-in algorithms containing knowledge about joins and data volume. It matches this intelligence against the join predicates, databases statistics, and available indexes to estimate which order is more efficient. In general, the optimizer will deploy an algorithm that minimizes the number of times the inner table must be accessed for qualifying outer table rows. However, none of today’s relational optimizers are perfect.

Access Path Choices

The relational optimizer has numerous options for creating SQL access paths. We have already discussed the join access paths, so let’s discuss some of the other common types of data access.

Table Scans

Table scans are the simplest form of data access. A table scan is performed simply by reading every row of the table. Depending on the DBMS, an alternate type of scan may exist, called a *tablespace scan*. The *tablespace*

scan reads every page in the tablespace, which may contain more than one table. Obviously, a tablespace scan will run slower than a table scan because additional I/O will be incurred reading data that does not apply.

Table scans are the simplest form of data access.

Another form of scanning is the *partition scan*. If the DBMS can determine that the data to be accessed exists in certain partitions of a multipartition table (or tablespace), it can limit the data that is scanned to the appropriate partitions. A partition scan should outperform a table scan or tablespace scan because the amount of I/O required is reduced.

Typically, the optimizer will choose to scan data for one of the following reasons:

- The query cannot be satisfied using an index, possibly because no index is available, no predicate matches the index, or the predicate precludes the use of an index.
- A high percentage of the rows in the table qualify. In this case, using an index is likely to be less efficient because most of the data rows need to be read anyway.
- The indexes that have matching predicates have low cluster ratios and are efficient for only small amounts of data.
- The table is so small that use of an index would actually be detrimental. For small tables, adding index access to the table access can result in additional I/O, instead of less I/O.

To assist the performance of a scan, the optimizer can invoke data prefetch. Data prefetch causes the DBMS to read data pages sequentially into the data cache even before they are requested. Essentially, data prefetch is a read-ahead mechanism—when data scans get around to requesting the data, it will already exist in memory. Data prefetch is particularly useful for table and tablespace scans but can be practical for any type of sequential data access. Whether data prefetch is available, as well as when and how it is used, depends on the DBMS. The optimizer may choose to deploy it when the access path is formulated, or the DBMS may choose to turn on data prefetch when the query is being run. As a DBA, you should learn how and why your particular DBMS prefetches data.

Data prefetch is particularly useful for table and tablespace scans.

Indexed Access

Of the many decisions that must be made by the optimizer, one of the most important for query performance is whether an index will be used to satisfy the query. To determine this, the optimizer must first discover whether an index exists. An index does not have to be defined before SQL can be written to access a column—you can query any column of any table known to the database.

Of course, before the relational optimizer will use an index to satisfy a query, an appropriate index must already exist. Additionally, at least one indexed column must be referenced within an indexable predicate in the SQL statement. The DBMS is not capable of using an index for every WHERE clause. You must learn what types of predicates can use indexes to ensure that the appropriate indexes are created for the queries in your database applications. Every DBMS has a different list of what is, and what is not, indexable. Furthermore, what is indexable tends to change from version to version of each DBMS.

The relational optimizer can choose to use an index in many different ways. The first, and simplest, type of indexed access is the *direct index lookup*. To perform a direct index lookup, the DBMS initiates the following steps:

1. The value in the SQL predicate is compared to the values stored in the root page of the index. Based on this comparison, the DBMS will traverse the index to the next-lower set of pages.
2. If intermediate nonleaf pages exist, the appropriate nonleaf page is read, and the value is compared to determine which leaf page to access.
3. The appropriate leaf page is read; the index leaf page contains pointer(s) to the actual data for the qualifying rows.
4. Based on the pointer(s) in the leaf page index entries, the DBMS reads the appropriate table data pages.

The simplest type of indexed access is the direct index lookup.

For the DBMS to perform a direct index lookup, values must be provided

for each column in the index. For example, consider the following query:

[Click here to view code image](#)

```
SELECT last_name, first_name, middle_initial, empno
FROM employee
WHERE position = 'MANAGER'
AND work_code = 1
AND dept = '001000';
```

Furthermore, assume that an index exists on the position, work_code, and dept columns. The DBMS can perform a direct index lookup using the index and the values supplied in the predicate for each of the columns. For a direct index lookup to occur, all three columns must appear in the SQL statement. If only one or two of these columns are specified as predicates, a direct index lookup cannot be chosen because the DBMS cannot match the full index key. Instead, an index scan could be chosen. Recall the previous discussion of table and tablespace scans. Index scans are similar. When an index scan is invoked, the leaf pages of the index are read sequentially, one after the other.

There are two basic types of index scans: matching index scans and nonmatching index scans. A *matching index scan* is sometimes called *absolute positioning*. A matching index scan begins at the root page of an index and works down to a leaf page in much the same manner as a direct index lookup does. However, because the complete key of the index is not available, the DBMS must scan the leaf pages of the index looking for the values that are available, until all matching values have been retrieved. To see how a matching index scan works, follow the arrows in [Figure 12.3](#).

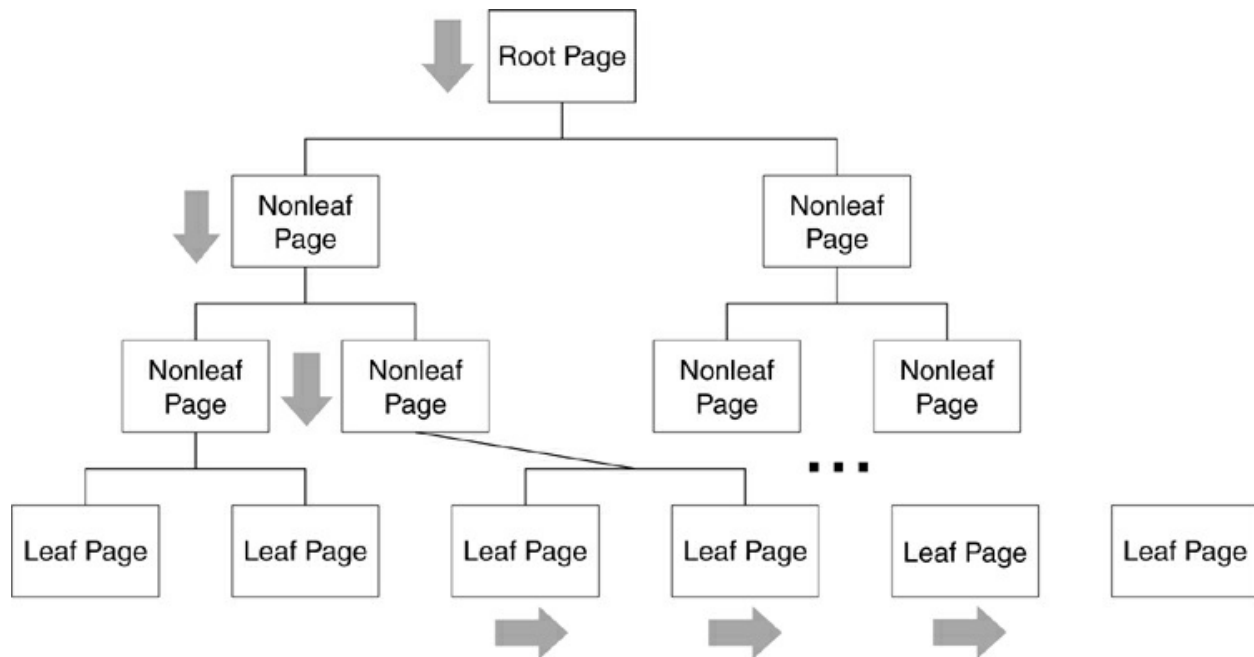


Figure 12.3. Matching index scan

There are two basic types of index scans: matching and nonmatching.

To clarify, consider the previous query of the employee table, but this time the query is recoded without the predicate for the dept column. The DBMS can use a matching index scan. The matching index scan first locates the first leaf page with the appropriate value for position and work_code by traversing the index, starting at the root. However, there can be multiple index entries with this combination of values and different values for dept. Therefore, the leaf pages will be scanned sequentially until no more valid entries are found.

For a matching index scan to be requested, you must specify the high-order column in the index key, in other words, for the first column specified in the index DDL. For the preceding example, the high-order column is the position column. The high-order column provides the starting point for the DBMS to traverse the index structure from the root page to the appropriate leaf page.

Consider the consequences of not specifying the high-order column in the query. For example, suppose we take the original query and remove the predicate for position but retain the other two, leaving the following SQL statement:

[Click here to view code image](#)

```

SELECT last_name, first_name, middle_initial, empno
FROM employee
WHERE work_code = 1
AND dept = '001000';

```

In such situations, the DBMS can deploy a *nonmatching index scan*, sometimes referred to as *relative positioning*. When a starting point cannot be determined because the first column in the index key is not specified, the DBMS cannot use the index tree structure. However, it can use the index leaf pages, as shown in [Figure 12.4](#). A nonmatching index scan begins with the first leaf page in the index and scans subsequent leaf pages sequentially, applying the available predicates.

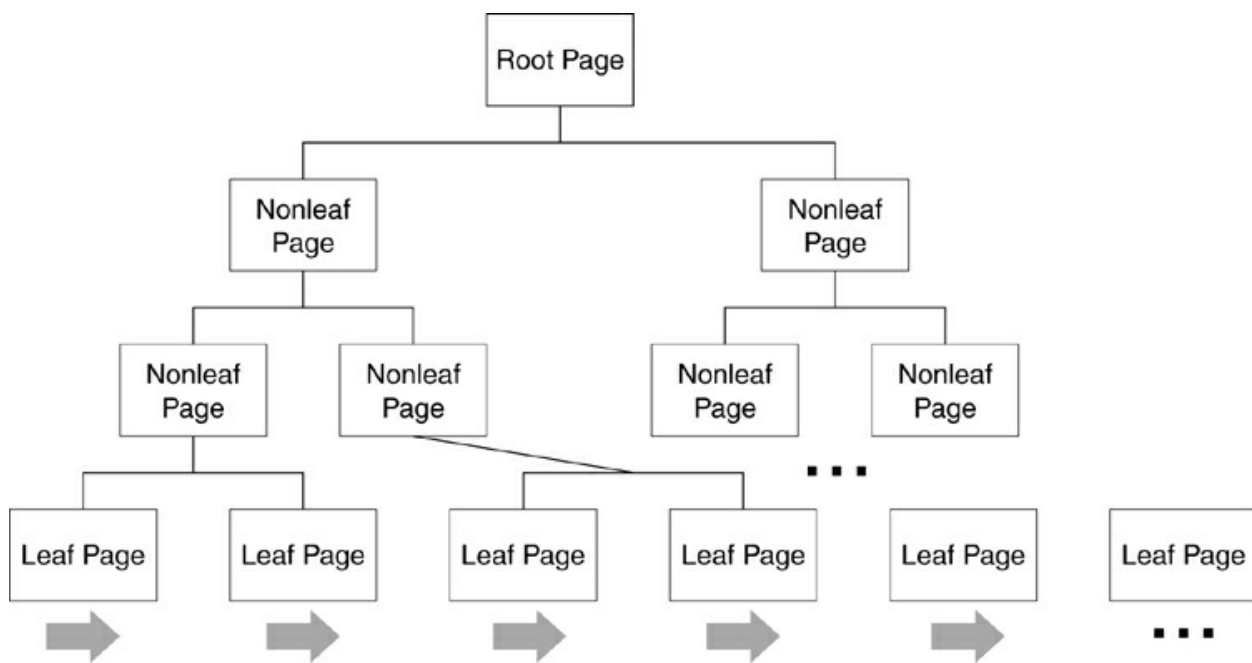


Figure 12.4. Nonmatching index scan

A nonmatching index scan can be more efficient than a table or tablespace scan, especially if the data pages that must be accessed are in clustered order. Of course, a nonmatching index scan can be done on a nonclustered index also.

This brings us to clustered versus unclustered index access. Any of the preceding methods for indexed access can be used with both clustered and unclustered indexes. However, index scans that must then access table data pages are not likely to be very efficient when the data is not clustered: Clustered index access, as it proceeds from leaf page to leaf page, never

requests a read for the same data page twice. Unclustered index access, however, will request the same data page multiple times because the data is spread throughout the table. For clarification of clustering, refer to the discussion and diagrams in [Chapters 4](#) and [11](#).

An additional index access technique that can be deployed is *index screening*. With index screening, a matching index scan is done on the leading columns of a multicolumn index, and additional predicates are applied during the scan. This technique is useful if some columns of a multicolumn index are not specified in the query. Consider another sample query:

```
SELECT last_name, empno
FROM employee
WHERE position = 'MANAGER'
AND work_code = 1
AND salary > 50000.00;
```

Index screening is useful if some columns of a multicolumn index are not specified in the query.

Now assume that an index has been created on the following columns in the following order: position, work_code, dept, and salary. The index can be screened by applying a matching index scan on position and work_code, and then a nonmatching scan for the specified salary values greater than \$50,000.00, but only for those rows that match the “position = ‘MANAGER’” and “work_code = 1” predicates.

One of the most efficient types of indexed access is *index-only access*, sometimes referred to as *index covering*. Consider again the query we just examined. Further, assume that an index has been created on the following columns: position, work_code, dept, salary, last_name, and empno. In such a scenario, the DBMS can satisfy the query using only the index because all of the requested data in the SELECT list and the predicates exists in the index. No additional I/O to table data pages is required.

A covering index contains all of the data required to satisfy the query without requiring access to the table.

With index-only access, the DBMS can satisfy the query by simply scanning the leaf pages of the index. A nonmatching index-only scan can be

much faster than a tablespace or table scan because index entries are generally smaller than table rows, and consequently more data is read with each I/O.

To encourage index-only access, the DBA can overload the index by adding extra columns that appear in the SELECT list of SQL statements. Doing so can provide a large performance payback, but it comes at the cost of additional disk storage requirements because additional columns are being indexed.

One final type of indexed access is *multi-index access*. With multi-index access, the DBMS uses more than one index to satisfy a single access path. For example, consider another variation of our employee table query:

```
SELECT last_name, empno
FROM employee
WHERE position = 'MANAGER'
AND work_code = 1;
```

With multi-index access, the DBMS uses more than one index to satisfy a single access path.

This time we have two indexes: one on the position column and another on the work_code column. The query specifies two predicates, each of which is supported by a different index. Instead of choosing to use one index or the other, the relational optimizer can combine the two indexes to return the correct data efficiently.

There are two types of multi-index access, depending on whether the predicates are tied together using AND or OR. Some DBMSs support only the AND logical operator, whereas others support both AND and OR. In some cases the DBA can minimize the number of indexes created by deploying multiple single-column indexes, instead of multiple multicolumn indexes. The DBA should base this decision on the efficiency of the DBMS in using multi-index access and the performance requirements of the applications being deployed.

Using Indexes to Avoid Sorts

The DBMS may need to sort data to satisfy SQL requests. Sorting is quite cost prohibitive and should be avoided if possible. The DBA can use indexes to avoid sorts by creating them on the columns that need to be sorted. The

relational optimizer will try to use indexes to avoid sorts whenever possible. Sorting might occur when the following clauses are specified:

- **DISTINCT:** When this clause is specified, the DBMS requires every column of the resulting data to be in order so that duplicate rows can be removed from the results set.

Use indexes to avoid sorts by creating them on the columns that need to be sorted.

- **UNION:** This operation requires the columns in each SELECT list to be ordered because the results set can have no duplicate rows. If the DBMS supports INTERSECT and EXCEPT, the same consideration applies to those operations.
- **GROUP BY:** When this clause is specified, the DBMS requires data to be sorted by the specified columns in order to aggregate data.
- **ORDER BY:** When this clause is specified, the DBMS will ensure that the results set is sorted by the specified columns.

Consider the following SQL statement:

[Click here to view code image](#)

```
SELECT  last_name, first_name, middle_initial, empno,
        position
FROM    employee
WHERE   position in ('MANAGER', 'DIRECTOR', 'VICE
PRESIDENT')
ORDER BY last_name;
```

If an index exists on the last_name column, the query can use this index and avoid sorting. Using an index to avoid a sort trades off the additional CPU cost required to sort for the additional I/O cost required for indexed access. Of course, if the index is going to be used anyway, the choice is a no-brainer. Whether or not using an index is actually faster than scanning the data and sorting will depend on

- Number of qualifying rows
- Speed of the sort
- Index characteristics (e.g., clustered or nonclustered)

Additionally, when coding queries that specify the UNION operation, take

care to examine the requirements of your application. If there could be no duplicate rows, or if you do not care if duplicates exist in the results set, you can specify UNION ALL to avoid the sort for duplicate removal.

Why Wasn't the Index Chosen?

Situations sometimes arise where you think the optimizer should have chosen an index, but it didn't. Any number of reasons can cause the optimizer to avoid using an index. Consult the following checklist for ways to encourage index selection:

- Does the query specify a search argument? If no predicate uses a search argument, the optimizer cannot use an index to satisfy the query.
- Are you joining a large number of tables? The optimizer within some DBMSs may produce unpredictable query plan results when joining a large number of tables.
- Are statistics current? If large amounts of data have been inserted, updated, and/or deleted, database statistics should be recaptured to ensure that the optimizer has up-to-date information upon which to base its query plans.
- Are you using stored procedures? Sometimes the DBMS provides options whereby a stored procedure, once compiled, will not reformulate a query plan for subsequent executions. You may need to recompile or reoptimize the stored procedure to take advantage of up-to-date statistics, new indexes, or any other pertinent database changes.
- Are additional predicates needed? A different WHERE clause might possibly enable the optimizer to consider a different index.

Hashed Access

The optimizer will also consider using any existing [hashing](#) structures when formulating access paths. A *hash*, or hash function, is an algorithm that converts a key (one or more columns) into a small number, usually a storage location on disk. The values returned by a hash function are called hash values, hash sums, or simply hashes.

Refer to [Figure 12.5](#) for a depiction of hash functionality. Key values are processed by the hash algorithm (also known as a hash function, hash routine, or randomizer) and translated into a storage location. When data is inserted, the hash algorithm tells the DBMS where to physically store the data; when

data is accessed by the key, the algorithm tells the DBMS where to find the data.

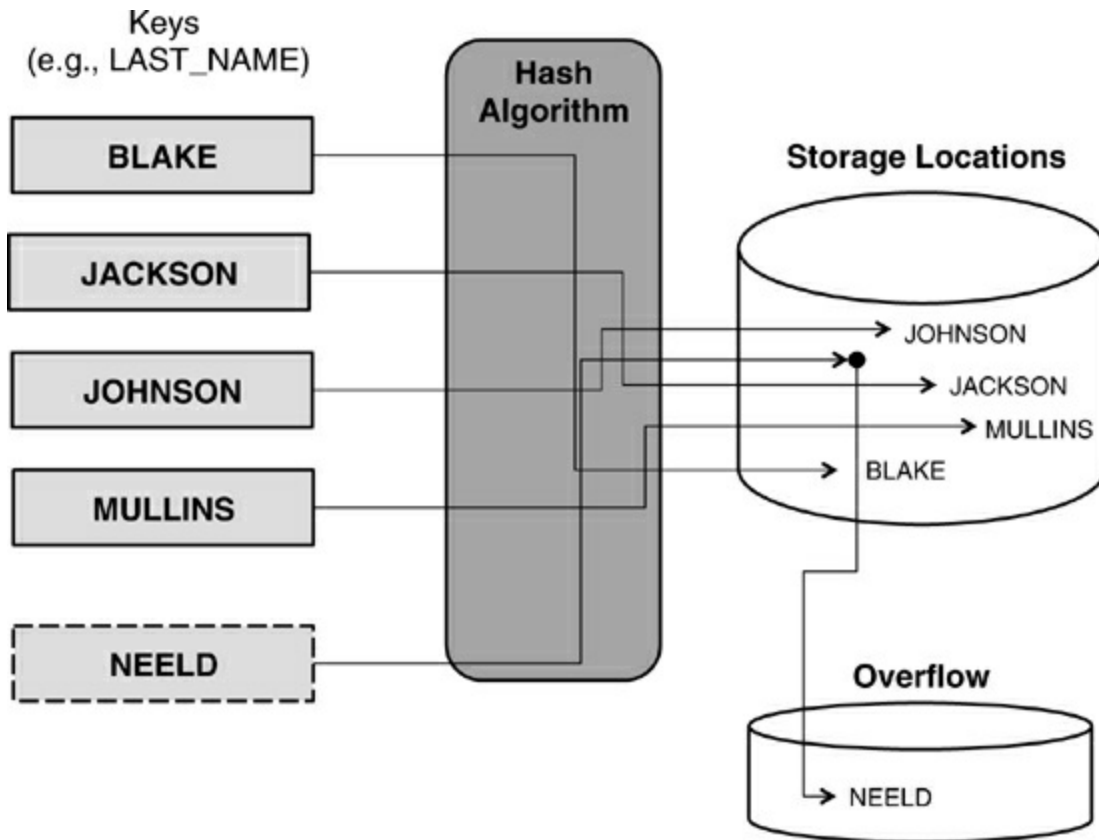


Figure 12.5. Hashing operation

A hash translates keys to storage locations.

A hash is more efficient than even a direct index lookup because it generally incurs less I/O. Instead of having to traverse through multiple index pages (from the root through nonleaf pages out to the leaf page and then to the data), a hash converts the key to a specific location on disk. This incurs a single I/O in the best-case scenario. If the hashing algorithm produces the same location for multiple keys, hash collisions occur, which require an additional I/O. A collision is shown when the key for NEELD is passed through the hash algorithm in [Figure 12.5](#).

Hashing is used primarily to optimize random I/O for small amounts of data, such as for looking up a code table value or accessing a single row based on the value of its primary key.

Parallel Access

The relational optimizer may choose to run queries in parallel. When query parallelism is invoked by the DBMS, multiple simultaneous tasks are invoked to access the data. Three basic types of parallelism can be supported by the DBMS:

- *I/O parallelism* enables concurrent I/O streams to be initiated for a single query. Running parallel I/O tasks can significantly enhance the performance of I/O bound queries. Breaking the data access for the query into concurrent I/O streams executed in parallel can reduce the overall elapsed time for the query.
- *CPU parallelism* enables multitasking of CPU processing within a query. Invoking CPU parallelism also invokes I/O parallelism because each CPU engine requires its own I/O stream. CPU parallelism decomposes a query into multiple smaller queries that can be executed concurrently on multiple processors. CPU parallelism can further reduce the elapsed time for a query.
- Finally, the DBMS can deploy *system parallelism* to further enhance parallel query operations. System parallelism enables a single query to be broken up and run across multiple DBMS instances. Allowing a single query to take advantage of the processing power of multiple DBMS instances can decrease the overall elapsed time for a complex query even further.

Three basic types of parallelism can be supported by the DBMS: I/O, CPU, and system.

Ensuring that proper query plans are formulated with the correct index usage is a time-consuming process, but one that can pay huge dividends in the form of enhanced performance. The DBA should train the application development staff to understand relational optimization and to create optimal SQL. Of course, the onus falls on the application developer to code efficient SQL and program logic. However, the DBA is the sentry of relational database performance. When performance problems occur, the DBA is the one who has to search for the cause of the problem and suggest remedies to resolve it. Furthermore, the DBA should conduct design reviews to seek out and tune inefficient SQL before suboptimal access paths and programs are migrated to production status.

Additional Optimization Considerations

The optimizer makes additional decisions regarding the manner in which data is accessed for SQL queries that will impact performance. In this section we will discuss some of these additional optimization considerations.

View Access

One of the decisions that must be made during query optimization is how to access data from views. Remember from [Chapter 4](#) that a view is a logical representation of a table that is defined using SQL. Therefore, a query that accesses a view is basically a SQL statement embedded within another SQL statement.

When the optimizer determines the access path for the query containing the view, it must also determine how to resolve the view SQL. Keep in mind that both the view and the SQL accessing the view may reference multiple tables and additional views.

Two methods can be used to optimize SQL that references views: view merging and view materialization. The more efficient of the two methods is *view merging*. As the name implies, when view merging is deployed, the SQL in the view DDL is merged with the SQL that references the view. The merged SQL is then used to formulate an access path against the base tables in the views.

View merging and view materialization can be used to optimize SQL that references views.

The second technique for accessing views is *view materialization*. When the optimizer cannot combine the SQL in the view with the SQL accessing the view, it creates an intermediate work file to hold the results of the view. The SQL accessing the view is then run against the work file that contains the view data. View materialization is not as efficient as view merging because data must be retrieved and stored in a temporary work file.

Each DBMS has its own set of rules that determine when view materialization must be used instead of view merging. Generally, column functions, or operations requiring sorts to be invoked, tend to require view materialization.

Query Rewrite

Some relational optimizers are intelligent enough to rewrite SQL more efficiently during the optimization process. For example, the optimizer might convert a subquery into an equivalent join. Alternatively, it might test out equivalent but different predicate formulations to determine which one creates the better access path. For example, since the following two predicates are equivalent, the optimizer may rewrite the query both ways to see which one produces the better access path:

[Click here to view code image](#)

```
WHERE column1 >= 1 AND column1 <= 100
WHERE column1 BETWEEN 1 AND 100
```

Additionally, the optimizer may rewrite queries by creating inferred predicates. One example of this is a feature known as *predicate transitive closure*, in which the optimizer adds a predicate to the query to improve performance. Consider the following SQL statement:

[Click here to view code image](#)

```
SELECT    d.dept_name, e.last_name, e.empno
FROM      employee   e,
          department d
WHERE     e.deptno = d.deptno
AND      d.deptno = '808';
```

That SQL statement is functionally equivalent to the following SQL statement:

[Click here to view code image](#)

```
SELECT    d.dept_name, e.last_name, e.empno
FROM      employee   e,
          department d
WHERE     e.deptno = d.deptno
AND      e.deptno = '808';
```

The only difference is the second predicate, but because deptno is the same in both tables (due to the first join predicate), it does not matter whether we check deptno from the employee table or the department table. However, it might make a difference in terms of performance. For example, an index might exist on one of the deptno columns but not the other, or perhaps one of the tables is significantly larger than the other. A query is usually more efficient when the predicate is applied to the larger of the two tables because

the number of qualifying rows will be reduced.

If the optimizer can perform predicate transitive closure, the SQL developer need not worry about this. The optimizer will consider the access path for both columns regardless of which is coded in the predicate. In essence, the optimizer will rewrite the query to add the redundant predicate.

The DBA should find out whether the relational optimizer in use can perform any form of query rewrite. Additionally, the rules for what types of queries can be rewritten vary from DBMS to DBMS and optimizer to optimizer. For example, the optimizer may not be able to perform predicate transitive closure on some predicates, such as IN or LIKE clauses.

The rules for what types of queries can be rewritten vary with DBMS and optimizer.

Rule-Based Optimization

The discussion up to this point has focused on cost-based optimization. Most relational optimizers are cost based, meaning they base their access path formulation decisions on an estimation of costs. Lower-cost access paths are favored over costlier access paths.

However, some DBMSs support a different type of optimization that is based on heuristics, or rules. Oracle, for example, provides both cost-based and rule-based optimization.

Oracle provides both cost-based and rule-based optimization.

A rule-based optimizer bases its optimization decisions on SQL syntax and structure, placement of predicates, order of tables in the SELECT statement, and availability of indexes. With a rule-based optimizer, SQL developers have to be aware of the rules as they write SQL. Query performance can suffer simply by reordering columns in the SELECT list or tables in the FROM clause.

Cost-based optimization is the trend for DBMSs because SQL statements need not be coded following a set of esoteric “rules.” Furthermore, an optimizer that estimates the cost of different access paths produces efficient query execution plans more reliably.

Reviewing Access Paths

The programmer or DBA can examine the access paths chosen by the relational optimizer. The commands and process used to accomplish this depend on the DBMS. Usually the command to externalize access paths is EXPLAIN or SHOWPLAN. Regardless of the name, the general process is depicted in [Figure 12.6](#).

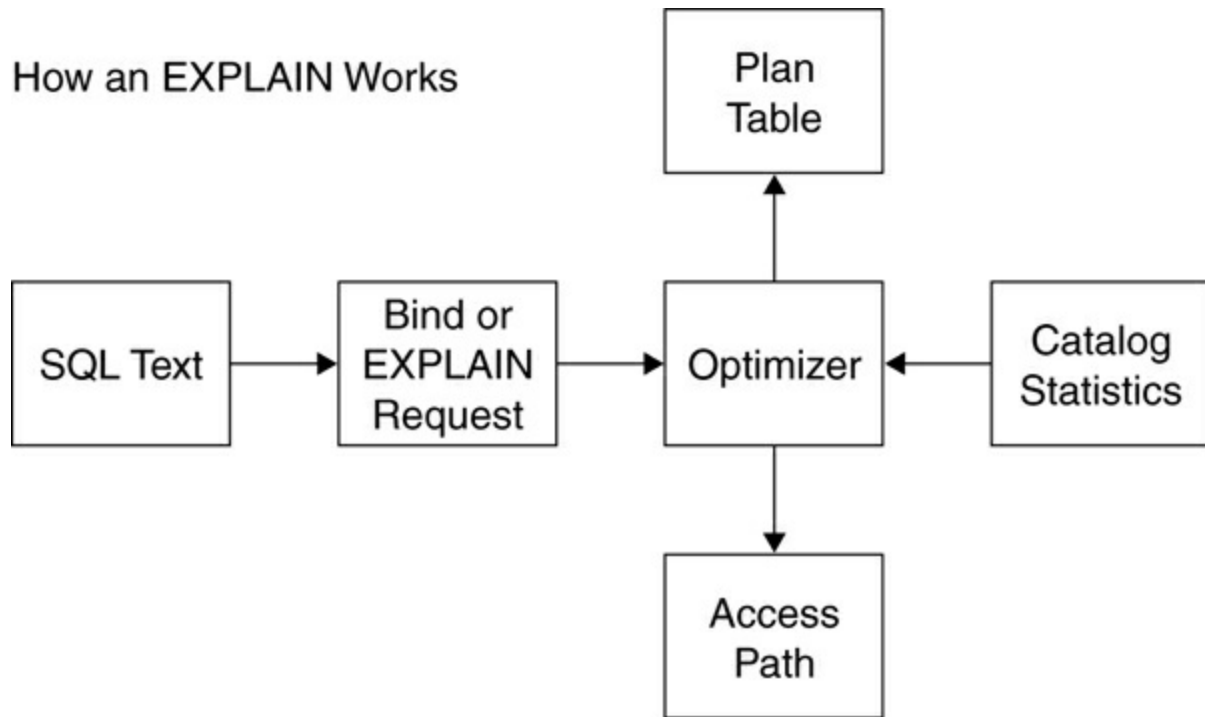


Figure 12.6. The EXPLAIN process

Microsoft SQL Server and Sybase make use of a command called SHOWPLAN. [Figure 12.7](#) provides a text description of the access path, and [Figure 12.8](#) provides a graphical depiction of the access path.

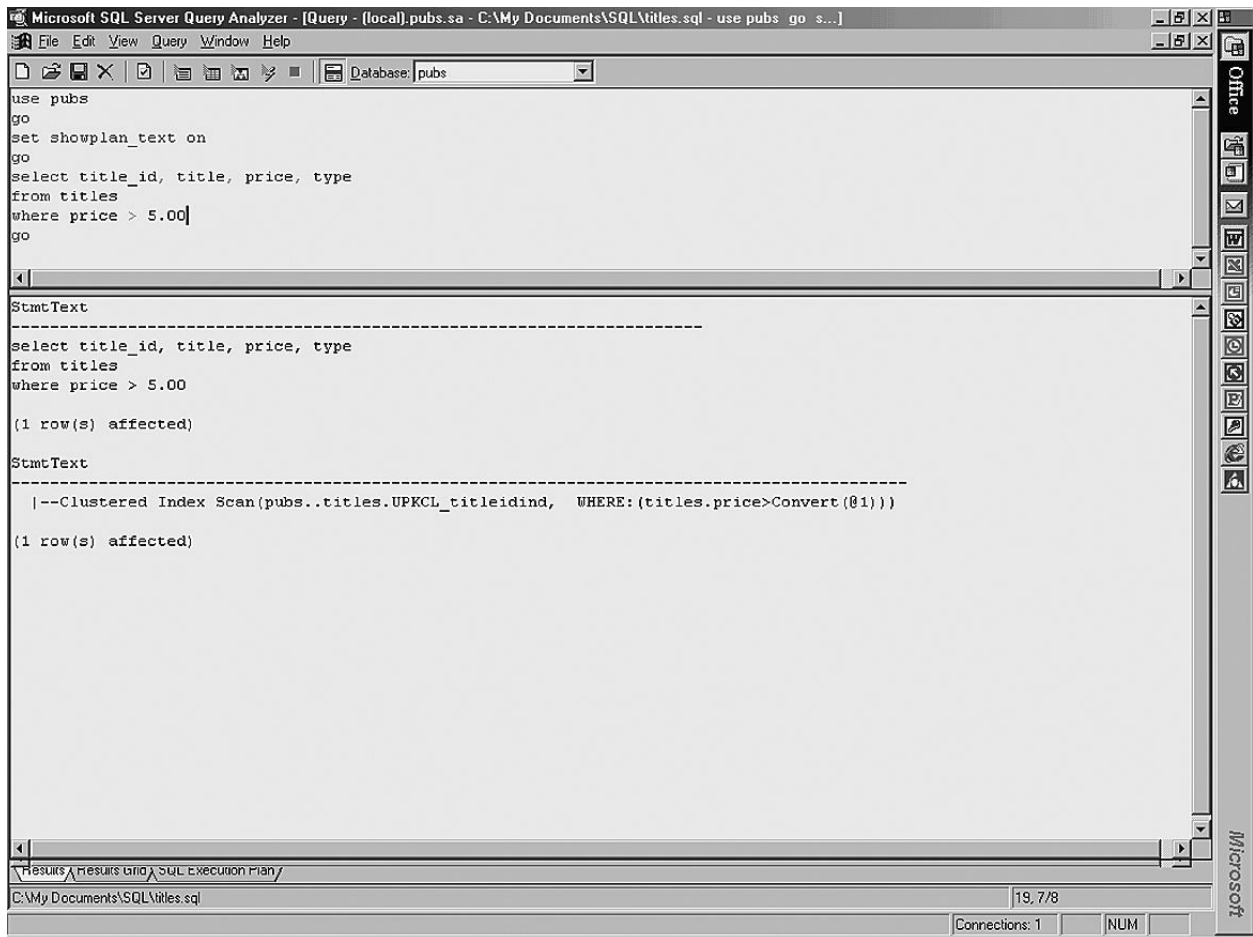


Figure 12.7. Textual SHOWPLAN example

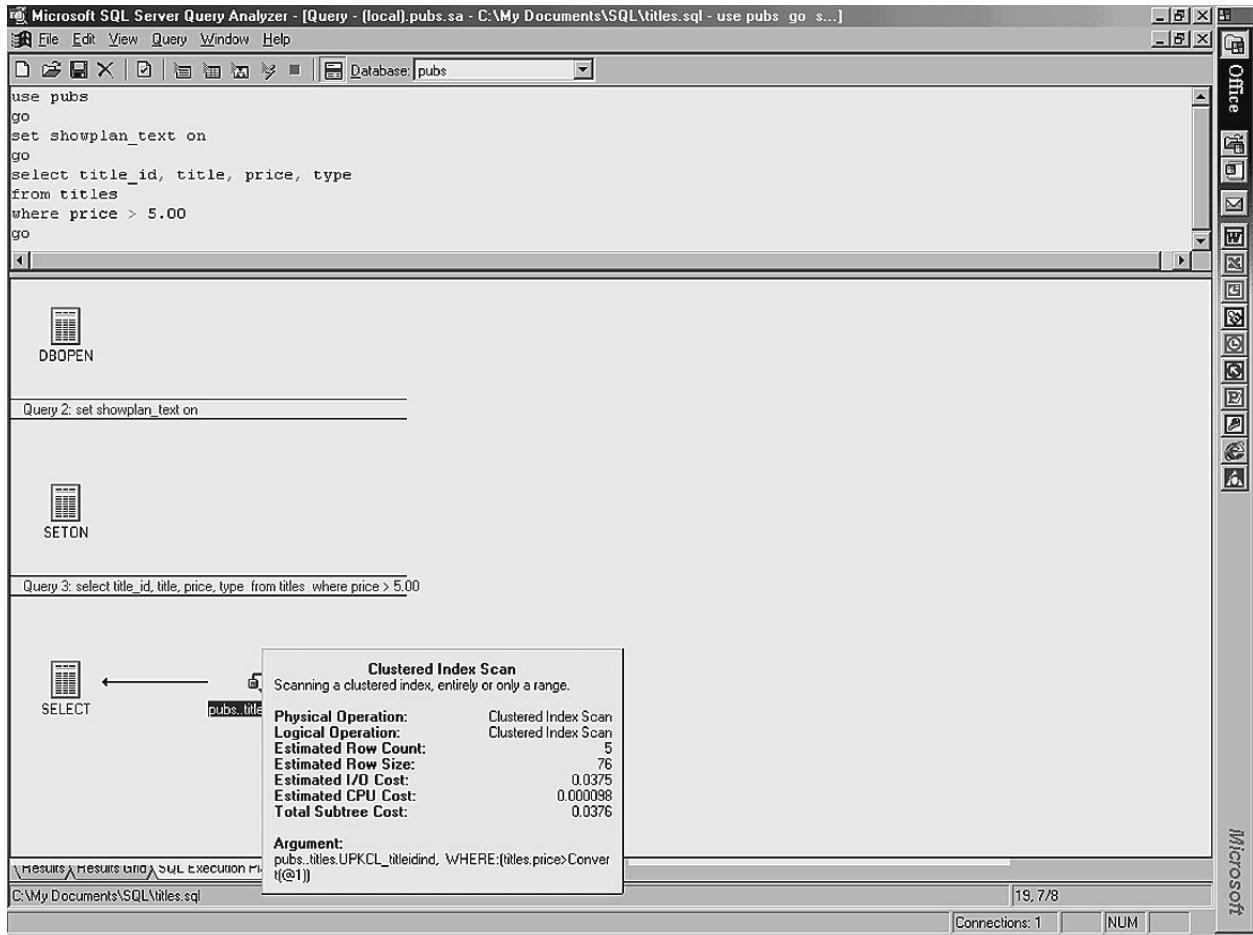


Figure 12.8. Graphical SHOWPLAN example

Oracle and DB2 use a different method, namely, the EXPLAIN statement. When the SQL statement is prefixed with the EXPLAIN command, the access path information determined by the optimizer is written out to a table called a PLAN_TABLE. The DBA or programmer can then query the PLAN_TABLE to interpret the access paths specified by the optimizer. EXPLAIN can be run on a statement-by-statement basis, or for an entire application program.

For example, consider the following SQL statement:

[Click here to view code image](#)

```
EXPLAIN plan SET STATEMENT_ID = 'emptest' FOR
SELECT position, last_name, first_name,
middle_initial, empno
FROM employee
WHERE position IN ('MANAGER', 'DIRECTOR', 'VICE
```



```
PRESIDENT')  
ORDER BY position;
```

Because the `SELECT` statement is prefixed with the `EXPLAIN` command, the relational optimizer will externalize access path information into the `PLAN_TABLE`. The `STATEMENT_ID` clause provides an identifier for locating the access path of the SQL statement within the `PLAN_TABLE`. The `PLAN_TABLE` contains encoded columns of data that outline the nature of query access paths to be used, such as the following:

- Whether an index is used, and if so, how many
- How many columns of the index match the query
- Whether index-only access is used
- What join method is used
- Whether parallel access is used
- Whether sorting is required

Of course, the exact columns and the permissible values of those columns differ between Oracle and DB2.

Access path analysis tools that interpret the `PLAN_TABLE` and display English text and graphical access path depictions are available for Oracle and DB2, too. For example, review the graphical output from the Visual Explain tool shown in [Figure 12.9](#). Such SQL analysis tools frequently provide expert tuning recommendations on how to fix inefficient SQL statements. Analysis tools can make it easier for the DBA to interpret the access paths being used and tune the SQL, if necessary.

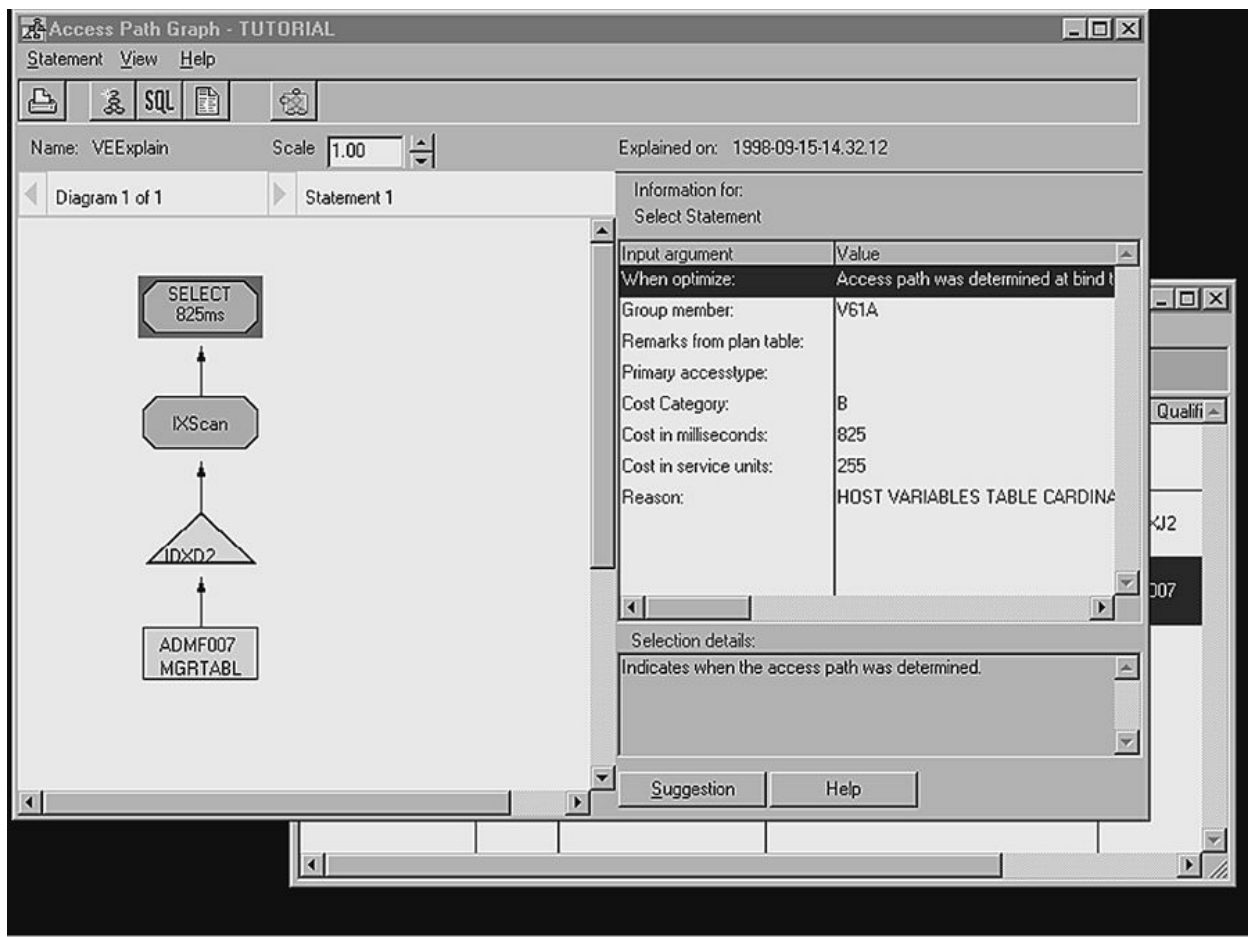


Figure 12.9. Visual Explain tool

Analysis tools can make it easier for the DBA to interpret the access paths being used.

Forcing Access Paths

Some DBMSs allow you to force the use of specific access paths or the order in which tables are joined. For example, Microsoft SQL Server provides the FORCEPLAN option. When FORCEPLAN has been set to ON, the optimizer will join the tables in the order in which they are coded in the SQL statement.

Oracle, on the other hand, provides hints that can be used to guide the relational optimizer to choose specific access paths. Hints are specified directly in the SQL query embedded within /*+ and */ , for example:

[Click here to view code image](#)

```

SELECT    /*+ USE_NL */
          e.position, e.last_name, e.empno, d.manager
FROM      employee  e,
          department d
WHERE     d.dept_id = e.dept_id
AND       position IN ('MANAGER', 'DIRECTOR', 'VICE
PRESIDENT')
ORDER BY position;

```

This query uses the Oracle hint `USE_NL` to force a nested-loop join. Additional hints can be provided to force the choice of index, force the use of parallelism, or force other optimization goals.

Techniques that force access path selection criteria should be used with caution. It is usually better to let the relational optimizer choose the appropriate access paths on its own unless

- You have in-depth knowledge of the amount and type of data stored in the tables to be joined
- You are reasonably sure that you can determine the optimal join order better than the optimizer, or
- Database statistics are not up-to-date, so the optimizer is not working with sufficient information about the database environment

Alternative methods are available to encourage the optimizer to select different access paths. The general method of encouraging access path selection is to modify the SQL, based on in-depth knowledge of the relational optimizer. This is sometimes called tweaking SQL. Since the optimizer within each DBMS is very different, few SQL tweaks are useful across multiple DBMSs.

Alternative methods are available to encourage the optimizer to select different access paths.

One common method of tweaking SQL is to change the query such that the results are the same, but the optimizer is incapable of using certain access paths. For example, consider the following SQL statement:

[Click here to view code image](#)

```

SELECT    last_name, first_name, empno, deptno
FROM      employee
WHERE     empno BETWEEN '001000' AND '009999'

```

```
AND      (salary > 50000.00 OR 0 = 1)
ORDER by last_name;
```

The results of this query are exactly the same with, or without, the “OR 0 = 1” component of the last predicate. However, some DBMS products prohibit the use of an index with such query formulations. In such cases, appending “OR 0 = 1” will eliminate the possibility of using an index to support the salary predicate.

DBAs must learn the fundamentals of SQL tuning and the types of tweaking that make sense for each DBMS that they manage. Furthermore, whenever such tweaks are deployed, be sure to fully document the reason. Without such documentation it is quite likely that a maintenance programmer will remove “OR 0 = 1” (or any other tweak) because it is not required.

The DBA must learn the fundamentals of SQL tuning.

SQL Coding and Tuning for Efficiency

Coding and tuning SQL is one of the most time-consuming DBA tasks. There can be literally thousands of individual SQL statements across hundreds of applications that access databases. The DBA is responsible for ensuring that the following steps occur for each SQL statement in the organization:

1. Identify the business data requirements.
2. Ensure that the required data is available within existing databases.
3. Translate the business requirements into SQL.
4. Test the SQL for accuracy and results.
5. Review the access paths for performance.
6. Tweak the SQL for better access paths.
7. Code optimization hints.
8. Repeat steps 4 through 7 until performance is acceptable.
9. Repeat step 8 whenever performance problems arise or a new DBMS version is installed.
10. Repeat the entire process whenever business needs change.

SQL tuning is a complex, time-consuming, and error-prone process, indeed. Furthermore, it requires cooperation and communication between the business users and application programmers for the first three steps, and

between the application programmers and the DBA for the remaining steps.

SQL tuning is a complex, time-consuming, and error-prone process.

Surely, there must be a way to simplify SQL tuning.

A Dozen SQL Rules of Thumb

In this section, we will examine some rules of thumb that apply generally to SQL development regardless of the underlying DBMS.

Rule 1: “It Depends!”

The answer to every question about database performance is “It depends.” A successful DBA will know on what it depends. For example, if someone asks, “What is the best access path for my SQL query?” the best answer is “It depends.” Why? Well, if every row must be returned, a table scan is likely to be more efficient than indexed access. However, if only one row is to be returned, direct index lookup will probably perform best. For queries that return between one and all rows, the performance of access paths will depend on how the data is clustered, which version of the DBMS is in use, whether parallelism can be invoked, and so forth.

A successful DBA will know on what it depends.

Be skeptical of tuning tips that use the words *always* or *never*. Just about everything depends on other things.

Rule 2: Be Careful What You Ask For

The arrangement of elements within a query can change query performance. To what degree depends on the DBMS in use and whether rule-based optimization is used.

A good rule of thumb, regardless of DBMS, is to place the most restrictive predicate where the optimizer can read it first. In Oracle, the optimizer reads WHERE clauses from the bottom up; therefore, the most restrictive predicate should be put at the bottom of the query. It is just the opposite in DB2.

Place the most restrictive predicate where the optimizer can read it first.

Placing the most restrictive predicate where the optimizer can read it first enables the optimizer to narrow down the first set of results before proceeding to the next predicate. The next predicate will be applied to the subset of data that was selected by the most selective condition, instead of against the entire table.

Rule 3: KISS

A rule of thumb for all types of IT activities is to follow the KISS principle: Keep it simple, stupid. However, in the world of SQL there is a trade-off between simplicity and performance.

Keep it simple, stupid.

Keeping SQL simple makes development and maintenance tasks easier. A simple SQL statement is easier to decipher and easier to change. With simple SQL, application developers can perform their job more easily than with complex SQL.

Nevertheless, complex SQL can outperform simple SQL. The more work that can be performed by the DBMS and the optimizer, the better performance is likely to be. Let's look at an example: Some programmers avoid joins by coding multiple SQL SELECT statements and joining the data using program logic. The SQL is simpler because the programmer need not understand how to write SQL to join tables. However, SQL joins usually outperform program joins because less data is returned to the program. Furthermore, the relational optimizer can change the join methodology automatically if the database or data changes. Conversely, program logic must be changed manually by a skilled programmer.

Rule 4: Retrieve Only What Is Needed

As simple as this rule of thumb sounds, you might be surprised at how often it is violated. To minimize the amount of data returned by your SQL statements, be sure to specify the absolute minimum number of columns in the SELECT list. If the column is not needed to satisfy the business requirement, do not request it to be returned in the results set.

Specify the absolute minimum number of columns in the SELECT list.

Programmers frequently copy SQL statements that work well to use as

templates for new statements. Sometimes programmers forget to trim down the number of columns requested when they need only a subset of the columns in the original query. This can adversely impact performance.

The more columns that must be returned by the DBMS, the greater the processing overhead.

Another common problem is requesting unnecessary data. Consider the following SQL statement:

```
SELECT    position, last_name, empno
FROM      employee
WHERE     last_name = 'SMITH';
```

There is no reason to specify the last_name column in the SELECT list of this SQL statement. We know that last_name must be SMITH for the entire results set because of the WHERE clause.

A similar type of problem can arise when programmers select a column only because they want to sort by it in the ORDER BY clause. Many DBMS products do not require the column to be selected in order to use it in an ORDER BY. The only columns that should be selected are those required for the business purpose driving the query.

Returning only what is needed does not apply only to columns. You should also minimize the number of rows to be returned by coding the proper WHERE clauses for every SQL statement. The more data that can be filtered out of the results set by the DBMS, the more efficient the query will be because less data must be returned to the requester.

Sometimes application programmers avoid coding appropriate WHERE clauses in a misguided attempt to simplify SQL statements. The more information the optimizer has about the data to be retrieved, the better the access paths it formulates will be. A sure sign of potential abuse is finding a SQL statement embedded in an application program that is immediately followed by a series of IF-THEN-ELSE statements. Try to tune the query by moving the IF-THEN-ELSE statements into SQL WHERE clauses.

Rule 5: Avoid Cartesian Products

Be sure to code predicates matching the columns of every table that will be joined within each SQL statement. Failure to do so will result in severe performance degradation and possibly incorrect results.

Whenever predicates do not exist for joining two tables, the RDBMS must perform a Cartesian product. This is the combination of *every* row of one table with *every* row of the other table. Nonmatching rows are not eliminated because there is nothing that can be matched. The results of a Cartesian product are difficult to interpret and contain no information other than a simple list of all rows of each table.

The results of a Cartesian product are difficult to interpret.

Rule 6: Judicious Use of OR

The OR logical operator can be troublesome for performance. If you can convert a SQL statement that uses OR to one that uses IN, it is likely that performance will improve. For example, consider changing this:

[Click here to view code image](#)

```
SELECT    e.position, e.last_name, e.empno, d.manager
FROM      employee   e,
          department d
WHERE     d.dept_id = e.dept_id
AND       position = 'MANAGER'
OR        position = 'DIRECTOR'
OR        position = 'VICE PRESIDENT'
ORDER BY position;
```

The OR logical operator can be troublesome for performance.

to this:

[Click here to view code image](#)

```
SELECT    e.position, e.last_name, e.empno, d.manager
FROM      employee   e,
          department d
WHERE     d.dept_id = e.dept_id
AND       position IN ('MANAGER', 'DIRECTOR', 'VICE
PRESIDENT')
ORDER BY position;
```

Of course, your results may vary depending on the DBMS in use and the nature of the data.

Rule 7: Judicious Use of LIKE

The LIKE logical operator is another troublesome beast. It is very easy to

create performance problems when using LIKE in SQL. For example, consider the following SQL:

[Click here to view code image](#)

```
SELECT    position, last_name, empno
FROM      employee
WHERE     dept_id LIKE '%X'
ORDER BY  position;
```

The LIKE logical operator is another troublesome beast.

This query will return employee information for all employees working in any department where dept_id ends in X. However, the relational optimizer will have to scan the data in order to resolve this query—there is no way to use an index. Because the high-order portion of the column is not known, traversing a b-tree index structure is impossible.

You might be able to use your knowledge of the data to rewrite this query without a leading wild-card character (%). For example, perhaps all dept_id values start with either A or B. In that case, you could modify the SQL as follows:

[Click here to view code image](#)

```
SELECT    position, last_name, empno
FROM      employee
WHERE     dept_id LIKE 'A%X'
OR        dept_id LIKE 'B%X'
ORDER BY  position;
```

In this case, the DBMS may be able to use a nonmatching index scan if an index exists on the dept_id column.

Once again, your results will vary with the DBMS in use and the nature of the data accessed.

Rule 8: Avoid Sorts When Possible

Sorting data is an inhibitor of optimal performance in SQL queries. Your DBMS will sort data as needed to satisfy your database requests. Follow the guidance in the “[Using Indexes to Avoid Sorts](#)” section earlier in this chapter for mitigating the impact of sorting when using ORDER BY, GROUP BY, DISTINCT, UNION, INTERSECT, and EXCEPT.

When performance is important, remember to look for sorts and find ways to eliminate them.

Keep in mind that sorting is an I/O-intensive operation and can degrade query performance, sometimes significantly. When performance is important, remember to look for sorts and find ways to eliminate them.

Rule 9: Know What Works Best

The flexibility of SQL allows the same process to be coded in multiple ways. However, one way of coding usually provides better performance than the others. The DBA should understand the best way to code SQL for each DBMS in use. Furthermore, the DBA should provide information on proper query formulation for performance to the application development staff.

One way of coding usually provides better performance than the others.

Rule 10: Issue Frequent COMMITs

When coding programs to run as batch transactions, it is important to issue regular SQL COMMIT statements. As discussed in [Chapter 5](#), the COMMIT statement finalizes modifications to the database. When a COMMIT is issued, locks on the modified database objects and data can be released.

An additional consideration for Oracle DBAs is the impact of a COMMIT on the rollback segments. Rollback segments are used by Oracle to store completed transactions before the changes are actually written to the table. When you issue a COMMIT in Oracle, not only is the data finalized to the table, but the contents of the rollback segment are removed, too. Oracle rollback segments are used to store “before” images of the data in case transactions are rolled back before changes are committed.

Therefore, as a DBA you must ensure that application developers issue enough COMMIT statements to minimize the impact of locking on availability and to keep rollback segments to a manageable size.

Issue enough COMMIT statements to minimize the impact of locking on availability.

Rule 11: Beware of Code Generators

Beware of application code generators and similar tools that automatically

create SQL. Code generators have a notorious reputation for creating poorly performing SQL. Be sure to keep a close watch on the SQL generated by such tools and use techniques to improve poorly written SQL before it reaches production status.

Code generators have a notorious reputation for creating poorly performing SQL.

Many of these tools use gateways that require each SQL statement to be recompiled and optimized each time it is requested. However, some gateways provide a caching mechanism to store compiled and optimized SQL on the server. Such a cache can help to improve performance for frequently recurring SQL statements.

Rule 12: Consider Stored Procedures

Performance degradation due to repeated network traffic can be minimized by using a stored procedure because only a single request is needed to execute it. Within the stored procedure, multiple SQL statements can be issued and the results processed and sent to the requesting program or user. Without the stored procedure, each of the multiple SQL statements, as well as all of the results, would have to be sent across the network. Additionally, SQL in stored procedures may perform better than the same SQL outside of the stored procedure if the DBMS parses and compiles the statements before run time.

SQL in stored procedures may perform better than the same SQL outside of the stored procedure.

Another potential benefit of using stored procedures is improving database security. This can be accomplished by encapsulating only the business logic needed for specific users into the code of the stored procedure and then granting access to the stored procedure instead of the underlying tables. This approach can diminish the potential security issues that ad hoc and dynamic SQL can cause (such as SQL injection attacks).

Additional SQL Tuning Tips

SQL tuning is a complicated task that requires a full-length book of its own (several good ones are referenced at the end of the chapter). In fact, each DBMS requires its own book-length treatment of SQL tuning to do the topic

justice. However, the following SQL tuning suggestions are useful for DBAs to apply, regardless of the DBMS:

- Create indexes to support troublesome queries.
- Whenever possible, do not perform arithmetic in SQL predicates. Use the host programming language (Java, COBOL, C, etc.) to perform arithmetic.
- Use SQL functions to reduce programming effort.
- Look for ways to perform as much work as possible using only SQL; optimized SQL typically outperforms host language application code.
- Build proper constraints into the database to minimize coding edit checks. (See [Chapter 13](#), “[Data Integrity](#),” for more details.)
- Do not forget about the “hidden” impact of triggers. A DELETE from one table may trigger many more operations. Although you may think the problem is a poorly performing DELETE, the trigger is really the culprit.

SQL tuning requires a full-length book of its own.

Identifying Poorly Performing SQL

A large part of the task of tuning SQL is identifying the offending code. A SQL performance monitor is the best approach to identifying poorly performing statements. Such a tool constantly monitors the DBMS environment and reports on the resources consumed by SQL statements. Some DBMSs provide rudimentary bundled support for SQL monitoring, but many third-party tools are available. These tools provide in-depth features such as the ability to identify the worst-performing SQL without the overhead of system traces, integration to SQL coding and tuning tools, and graphical performance charts and triggers. Consult [Chapter 23](#), “[DBA Tools](#),” for more information on SQL performance monitors and other DBA tools.

A SQL performance monitor is the best approach to identifying poorly performing statements.

Summary

Application performance management and SQL tuning is a complex area that requires the active participation of programmers and DBAs. Each DBMS

operates differently, and DBAs as well as programmers will need to understand all of the minute details of SQL and application performance management for their DBMS.

The relational optimizer combines access path strategies to form an efficient access path for each SQL request. However, the optimizer is a very complex piece of software, and DBMS vendors do not share with their customers all the intricacies of how their optimizer works. Therefore, quite often, SQL performance tuning becomes an iterative artistic process, instead of a science.

Review

1. Define what is meant by *physical data independence*.
2. Describe the impact of using the LIKE operation with a wild-card character at the beginning of the value.
3. What factors influence the optimizer during SQL access path selection?
4. Name and describe the two predominant relational join methods.
5. Under what circumstances is a nonmatching index scan performed?
6. A query is written to access a single table. Furthermore, that query will return only a single row because an equality operator (=) is coded on the primary key for the table. A unique index exists to support the primary key. What type of access is likely to be the most efficient for that query?
7. How can stored procedures be used to optimize performance in a client/server application?
8. Under what circumstances will a table scan outperform indexed access?
9. Why is it important to use caution before implementing SQL tweaks and hints?
10. Describe the three possible types of parallel query access.

Suggested Reading

Cumming, Andrew, and Gordon Russell. *SQL Hacks: Tips & Tools for Digging into Your Data*. Sebastopol, CA: O'Reilly (2007). ISBN 0-596-52799-3.

Faroult, Stephane. *The Art of SQL*. Sebastopol, CA: O'Reilly (2006). ISBN

0-596-00894-5.

Gulutzan, Peter, and Trudy Pelzer. *SQL Performance Tuning*. Boston, MA: Addison-Wesley (2003). ISBN 0-201-79169-2.

Harrington, Jan L. *SQL Clearly Explained*. 3rd ed. Burlington, MA: Morgan Kaufmann (2010). ISBN 978-0-12-375607-8.

Khoshafian, Setrag, et al. *A Guide to Developing Client/Server SQL Applications*. San Mateo, CA: Morgan Kaufmann (1992). ISBN 1-55860-147-3.

Lewis, Jonathan. *Cost-Based Oracle Fundamentals*. New York, NY: Apress (2006). ISBN 1-59059-636-6.

Molinaro, Anthony. *SQL Cookbook*. Sebastopol, CA: O'Reilly (2006). ISBN 978-0-596-00976-2.

Rozenshtein, David, et al. *Optimizing Transact-SQL: Advanced Programming Techniques*. Fremont, CA: SQL Forum Press (1995). ISBN 0-9649812-0-3.

Smirnova, Sveta. *MySQL Troubleshooting*. Sebastopol, CA: O'Reilly (2012). ISBN 978-1-449-31200-8.

Tow, Dan. *SQL Tuning*. Sebastopol, CA: O'Reilly (2004). ISBN 0-596-00573-3.

Williams, Hugh E., and David Lane. *Web Database Applications with PHP and MySQL*. 2nd ed. Sebastopol, CA: O'Reilly (2004). ISBN 0-596-00543-1.

13. Data Integrity

Ensuring the integrity of the organization's databases is a key component of the DBA's job. A database is of little use if the data it contains is inaccurate or if it cannot be accessed due to integrity problems. DBAs have many tools at their disposal to ensure data integrity.

DBAs have many tools at their disposal to ensure data integrity.

Types of Integrity

With respect to databases, we will discuss two aspects of integrity:

- Database structure integrity
- Semantic data integrity

Keeping track of database objects and ensuring that each object is created, formatted, and maintained properly is the goal of database structure integrity. Each DBMS uses its own internal format and structure to support the databases, tablespaces, tables, and indexes under its control. System and application errors can cause faults within these internal structures, and it is the DBA's job to identify and correct such faults before insurmountable problems are encountered.

The second type of integrity is semantic data integrity. This refers to the meaning of data and the relationships that need to be maintained between different types of data. The DBMS provides options, controls, and procedures to define and assure the semantic integrity of the data stored within its databases. DBAs must understand how the DBMS enables automatic semantic data integrity checking. And, as an ongoing component of the job, the DBA has to implement semantic data integrity into the database design, as well as initiate processes to check and correct integrity problems that creep into the database over time.

Database Structure Integrity

Structural database integrity and consistency are of paramount importance in the ongoing administration of databases. The DBMS uses internal structures

and pointers to maintain database objects in the proper order. If those structures are damaged in any way, database access will be compromised.

The DBMS uses internal structures and pointers to maintain database objects in the proper order.

Types of Structural Problems

One potential problem experienced by relational databases is index corruption. You will recall from previous chapters that an index provides an alternate path to data in the database by using an ordered b-tree structure. Basically, the leaf pages of an index are pointers to physical data locations in the base table. If the pointers do not point to the correct data, the index is useless. In fact, it could be worse than useless if it is used to access incorrect data and causes transactions to corrupt the database.

If the pointers do not point to the correct data, the index is useless.

An index can become corrupted in several ways, depending on the DBMS in use and the operations performed on the index and the corresponding table. For example, if the table is recovered to a previous point in time but the index is not rebuilt, the index will not accurately reflect the current contents of the data.

However, indexes are not the only database objects that use pointers. When certain data types are used, the DBMS may require pointers. For example, *very large objects* such as DB2 LOB columns and XML columns, or SQL Server text and image columns, are not stored contiguously with the rest of the data. Instead, the columns assigned these data types are stored in a separate file, and the primary file for the table contains a pointer to the physical data location of the large object. It is possible for these pointers to get out of sync with the data, rendering the data inaccessible.

Another, though somewhat rare, database structural integrity problem is *table data page header corruption*. Each individual physical page (or block) in the database file stores housekeeping information, known as the *header*, at the beginning of the page. This header enables the DBMS to easily and rapidly read the page contents. If the header becomes corrupted, the DBMS may not be able to interpret the data stored on the page. Such situations

usually require the database file to be recovered from backup files. (For more information on backup and recovery operations, refer to [Chapter 16](#).)

Backup files are another potential area for a structural integrity problem. Each DBMS formats and maintains a specific structure for its database backup files. If the backup file is not formatted correctly, or if data is in the wrong location in the backup file, it cannot be used by the DBMS for recovery purposes. Media failures, tape degradation, and bugs can cause such problems.

Managing Structural Problems

Each DBMS provides different utility programs to check on different aspects of database structural integrity. It is possible to investigate the integrity of a database using these utility programs. For example, Sybase and Microsoft SQL Server provide the DBCC utility program, DB2 supplies the CHECK and REPAIR utilities, and Informix offers the TBCHECK utility.

You can investigate the integrity of a database using DBMS utility programs.

To elaborate on the functionality of such utilities, let's focus on the DBCC utility. [Figure 13.1](#) shows some of the options for which DBCC can be used. Another example of a structural integrity utility is DB2's REPAIR utility, which is briefly discussed in the sidebar titled "[The DB2 REPAIR Utility](#)."

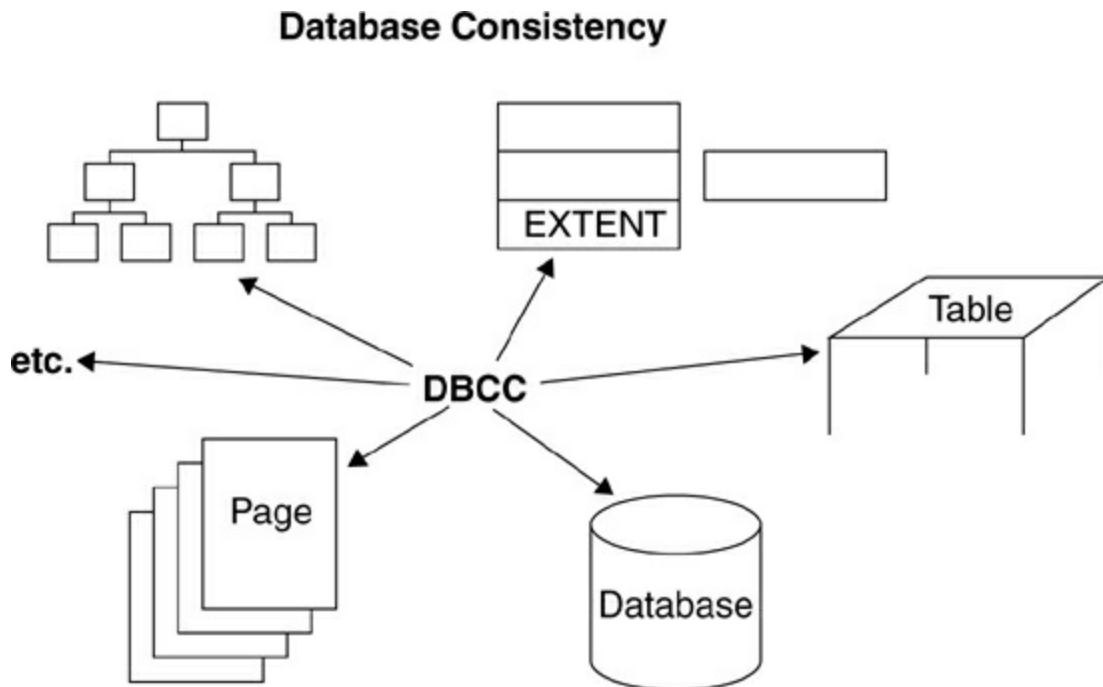


Figure 13.1. The DBCC utility

The DBA should exhibit extreme care when running the DBCC utility. It is possible to both read and write to database files when using DBCC. To ensure the safety of your database files, never use the write capabilities without guidance from the DBMS vendor's technical support group.

Use extreme care when running the DBCC utility.

Consistency Options

The DBCC utility provides two options for basic consistency checking:

- `DBCC CHECKTABLE(table_name)` checks the consistency of the data and index pages of a table. When DBCC is run using this option, it will report on the number of data pages, the number of rows, the number of text and image columns, as well as any integrity violations.

The DB2 REPAIR Utility

The REPAIR utility is designed to modify DB2 data and associated data structures when there is an error or problem. You can use the REPAIR utility to perform tasks such as testing DBD definitions, repairing DBDs by synchronizing with the system catalog metadata, resetting pending states, producing dumps, and surgically replacing specific data areas of

tablespaces and indexes (among other functions).

As with all database structural integrity utilities, care should be taken when using the REPAIR utility. Furthermore, it is wise first to consult with your DBMS support center (in this case, IBM) before running such utilities. Certain options of the REPAIR utility can damage your database and the data stored therein if used improperly.

- DBCC REINDEX(table_name) defragments the indexes built on the specified table.

Database Checking

Additional DBCC options exist for checking the integrity of a database:

- DBCC CHECKDB(database_name) runs CHECKTABLE (see above) on every table in the database. This option checks the consistency of the data and index pages of all defined tables.
- DBCC CHECKCATALOG(database_name) checks the consistency of the system catalog tables for the specified database. It will report on the size and number of segments used, as well as detect and report on any integrity errors.
- DBCC CHECKALLOC(database_name) checks the consistency of the specified database and reports on the current extent structure. The CHECKALLOC option may find spurious allocation problems if DBCC is run while transactions are modifying the database. This option also reports on the number of allocations and the pages used per allocation.
- DBCC CHECKFILEGROUP(filegroup_name) checks the allocation and structural integrity of all tables and indexed views in the specified database and issues a report of the findings.

Memory Usage

The DBCC command can be used to monitor current memory allocation and usage. If the MEMUSAGE keyword is specified, for example, DBCC MEMUSAGE, the utility will report on the configured memory allocation and memory usage of the top 20 memory users. The following information is displayed:

- *Configured memory*—the amount of memory configured for the DBMS

- *Code size*—the amount of memory used by code
- *Kernel and structures*—the amount of memory used for the kernel and the server

The DBCC command can be used to monitor current memory allocation and usage.

- *Page cache*—the amount of memory used for data cache
- *Procedure buffers and headers*—the amount of memory used for procedure cache
- *Buffer cache detail*—a listing of the top 20 users of data cache, indicating the database and object being accessed and how many 2K buffers are being used
- *Procedure cache detail*—a listing of the top 20 users of procedure cache, indicating the type of usage and the memory utilization

Additional Options

Various other DBCC options are available, including the ability to generate reports containing information on database internals (for example, creation date and internal identifier), to print formatted table pages showing the header contents, to dump and display the contents of buffers and caches, as well as to “zap” the database (that is, make quick changes to any of the database contents). Once again, though, it is not a good idea to “zap” the database without guidance from the manufacturer.

Semantic Data Integrity

Semantic data integrity deals with the DBMS features and processes that can be used to ensure the accuracy and viability of the data contents of a database. Whereas structural database integrity refers to the consistency of the “holders” of the data (the database objects), semantic data integrity refers to the consistency of the data itself.

DBAs constantly struggle with the question of how best to enforce data integrity—by using DBMS features or by using application code. This is an important issue and is debated often. In general, using DBMS features to support data integrity usually offers the better solution, for several reasons.

It is easier to discover and maintain a semantic integrity rule by using

DBMS features rather than by reading program code. If you prefer to read code, remember that you may need to be able to read several languages (COBOL, Java, Visual Basic, etc.) in a heterogeneous environment, often for the same rule. Additionally, you will need to know how to distinguish the active code from the old code from the new enhancement code.

Many forms of semantic data integrity can be enforced by using features of the DBMS. Refer to the sidebar “[SQL Server 2012 Semantic Integrity](#)” for a quick overview of SQL Server’s support in this area. When DBMS facilities are used to enforce data integrity, less code needs to be written, and therefore less code needs to be tested. This can reduce the time and cost to get the “product” out the door.

Many forms of semantic data integrity can be enforced by using features of the DBMS.

SQL Server 2012 Semantic Integrity

Microsoft SQL Server 2012 includes several features for managing semantic integrity and data quality. Perhaps the most intriguing featuring is Data Quality Services (DQS), which facilitates the ability for data professionals to maintain the quality of their data and ensure that the data is suited for the business purpose. DQS is a knowledge-driven solution that provides both computer-assisted and interactive ways to manage the integrity and quality of SQL Server data. DQS enables you to discover, build, and manage knowledge about your data. You can then use that knowledge to perform data cleansing, matching, and profiling. You can also leverage the cloud-based services of reference data providers in a DQS data quality project.

DQS provides the following features to resolve data quality issues:

- Data cleansing
- Data matching
- Reference data services
- Profiling
- Monitoring
- Knowledge base

SQL Server also features Master Data Services for the management of

master data. Using Microsoft SQL Server Master Data Services, a data professional can ensure the integrity of information and consistency of data across different applications. The Master Data Services software helps to ensure quality in a master data hub using rules and automated workflows.

Before investigating third-party software solutions for semantic data integrity, be sure you thoroughly investigate the functionality and feature set of your in-house DBMS products. Many contain features that can be used to improve the quality and integrity of your database data.

Entity Integrity

Entity integrity is the most basic level of data integrity provided by relational databases. [Entity integrity](#) means that each occurrence of an entity must be uniquely identifiable. In other words, entity integrity requires the specification of a primary key for each entity and that no component of the primary key be set to null.

Entity integrity means that each occurrence of an entity must be uniquely identifiable.

In practice, no major RDBMS product forces entity integrity because entities, or tables, can be created without specifying a primary key. However, it is considered a bad practice to create tables without a primary key because it makes identifying rows of the table difficult.

A primary key is required to set up referential integrity (RI) between tables. A primary key constraint can consist of one or more columns from the same table that are unique within the table. A table can have only one primary key constraint, which cannot contain nulls.

The DBMS will require the creation of a unique index on the columns of the primary key constraint to prohibit duplicate column values. Some DBMSs, such as Sybase, will automatically generate the unique index when the primary key constraint is created.

An example of a primary key constraint follows:

[Click here to view code image](#)

```
CREATE TABLE EMP
```

```

(empno          INTEGER          PRIMARY KEY,
 emp_address    VARCHAR(70),
 emp_type       CHAR(8),
 emp_dept       CHAR(3)          NOT NULL WITH DEFAULT,
 salary         DECIMAL(7,2)     NOT NULL,
 commission     DECIMAL(7,2),
 bonus          DECIMAL(7,2)
) IN db.ts;

```

This example shows the primary key constraint specified at the column level: It applies to the empno column, which is defined as the primary key for this table. However, in practice it is quite common for primary keys to consist of multiple columns. For this reason, constraints can be defined at the table level, as well. For example, if the primary key of this table were to be defined as the combination of the empno and emp_type columns, you could add the following specification to the end of the table, right after the specification of the last column defined:

```

PRIMARY KEY pkemp
(empno, emp_type)

```

Of course, the primary key specification would have to be removed from empno at the column level before this would work because you cannot specify more than one primary key per table. Furthermore, notice that we have named the primary key constraint pkemp in this example.

Unique Constraints

A *unique constraint* is similar to a primary key constraint. Each table can have zero, one, or many unique constraints consisting of one or more columns each. The values stored in the column, or combination of columns, must be unique within the table—that is, no other row can contain the same value.

Unique constraints differ from primary key constraints in that they cannot be used to support referential constraints. Furthermore, the unique constraint columns may be set to null.

Unique constraints cannot be used to support referential constraints.

As in the case of primary key constraints, the DBMS will require the creation of a unique index on the columns of the unique constraint to prohibit

duplicate column values. The DBA should create unique constraints for columns, or combinations of columns, that must be unique within the table. Unique constraints are more efficient than trying to enforce uniqueness programmatically.

Data Types

Data type and *data length* are the most fundamental integrity constraints applied to data in the database. Simply specifying the data type for each column when a table is created will enable the DBMS to automatically ensure that only the correct type of data is stored in that column. Processes that attempt to insert or update the data to the wrong type will be rejected. Furthermore, a maximum length is assigned to the column to prohibit larger values from being stored in the table.

The DBA must choose the data type and length of each column wisely. It is best to choose the data type that most closely matches the domain of correct values for the column. For example, a numeric column should be defined as one of the numeric data types: integer, decimal, or floating-point.

The DBA must choose the data type and length of each column wisely.

Specifying a character data type for a column that will contain numeric data imposes additional work on the application programs to edit-check the values. Furthermore, the program is likely to be less efficient at edit-checking data than the DBMS would be. Additionally, data that is inserted or modified outside the scope of an application program will not be edit checked at all and thereby has the potential of introducing data integrity problems into the database.

User-Defined Data Types

A *user-defined data type*, or UDT for short, provides a mechanism for extending the type of data that can be stored in databases and the way that the data is treated. In other words, the DBA can create user-defined data types to further clarify the legal values for a column. The UDT, once defined and implemented, extends the functionality of the DBMS by enabling the DBA to specify the UDT in CREATE TABLE DDL statements just like built-in data types.

A user-defined data type extends the type of data that can be stored in databases.

UDTs can be beneficial when you need to store data that is specifically tailored to your organization's requirements. For example, UDTs are useful for organizations that must handle multiple currencies. When there is a separate data type for each currency, errors are less likely to be made in monetary calculations. For example, perhaps your company handles monetary amounts from Canada, the United States, the European Union, and Japan. The DBA can create four UDTs, such as

[Click here to view code image](#)

```
CREATE DISTINCT TYPE canadian_dollar AS DECIMAL(11,2);
CREATE DISTINCT TYPE US_dollar AS DECIMAL(11,2);
CREATE DISTINCT TYPE euro AS DECIMAL(11,2);
CREATE DISTINCT TYPE japanese_yen AS DECIMAL(15,2);
```

After a user-defined data type has been created, it can be used in the same manner as a system-defined data type. Of course, the actual syntax for creating and using UDTs differs from DBMS to DBMS. Sybase and Microsoft SQL Server use system-stored procedures to create UDTs, whereas Oracle and DB2 use CREATE statements such as those in the previous example.

When separate data types are created, the DBMS enforces strong typing; in other words, the DBMS will prohibit nondefined operations between different types. For example, the following operation will not be allowed due to strong typing:

```
TOTAL_AMT = US_DOLLAR + CANADIAN_DOLLAR
```

If this operation were allowed to proceed, the TOTAL_AMT calculated would be useless because you cannot perform operations on different currencies without converting them to a common currency. So, UDTs and strong typing help us to avoid data integrity problems.

Another benefit of creating user-defined data types is additional consistency throughout the database. For example, if we build a Social Security number UDT, and use the UDT for every column that contains Social Security numbers, we ensure that each Social Security number is stored the same way.

Furthermore, UDTs provide a higher level of abstraction to your database design. It is much easier to discuss the Salary data type (with all its implied definitions, properties, and constraints) than it is to talk about DECIMAL(11,2) or SMALLMONEY data types (with no implied characteristics other than their inherent type).

Default Values

When columns are created within tables, they can be assigned a default value that will be used when SQL INSERT statements are issued that do not provide an explicit value for that column. This allows programmers to ignore columns and the DBMS to automatically provide a default value.

Each column can have only one default value. The column's data type, length, and property must be able to support the default value specified. For example, a numeric column cannot be assigned a character default value. Furthermore, the default may be null, but only if the column is created as a nullable column.

Each column can have only one default value.

Check Constraints

A *check constraint* is a DBMS-defined restriction placed on the data values that can be stored in a column or columns of a table. When a check constraint is defined, it places specific data value restrictions on the contents of a column through the specification of a Boolean expression. The expression is explicitly defined in the table DDL and is formulated in much the same way that SQL WHERE clauses are formulated. Any attempt to modify the column data (that is, during INSERT and UPDATE processing) will cause the expression to be evaluated. If the modification conforms to the Boolean expression, the modification is permitted to proceed. If not, the statement will fail with a constraint violation.

A check constraint is a DBMS-defined restriction placed on the data values that can be stored in a column or columns.

Check constraints can be defined when the table is created, or added later by altering the table. If you know SQL, it is easy to learn how to code check constraints because they are written using recognizable SQL syntax. The check constraint syntax consists of two components: a constraint name and a

check condition.

The *constraint name* identifies the check constraint to the database. The same constraint name cannot be specified more than once for the same table. If a constraint name is not explicitly coded, the DBMS automatically generates a unique name for the constraint. Each DBMS uses a different algorithm for generating constraint names, but the name is typically derived from the name of the first column in the check condition.

The *check condition* defines the actual constraint logic. The check condition can be defined using any of the basic predicates (>, <, =, <>, <=, >=), as well as BETWEEN, IN, LIKE, and NULL. Furthermore, AND and OR can be used to string conditions together in a check constraint.

Some restrictions, however, apply to creating check constraints:

- The check constraint can refer only to columns in the table in which it is created.
- A limited subset of SQL constructs is permissible within a check constraint definition. Typically, SQL constructs such as subselects, column functions, host variables, negation (i.e., NOT), and special registers are prohibited within check constraints.
- The first operand of the check constraint is the name of a column contained in the table; the second operand is either another column name or a constant.
- If the second operand is a constant, it must be compatible with the data type of the first operand. If the second operand is a column, it must be the same data type as the first column specified.

Check Constraint Benefits

So, what are the benefits of check constraints? The primary benefit is the ability to enforce business rules directly into the database without requiring additional application logic. Once defined, the business rule is physically implemented and cannot be bypassed.

Because no additional programming is required, DBAs can implement check constraints without involving the application programming staff. This effectively minimizes the amount of code that must be written by the programming staff. Given the significant application backlog within most organizations, this can be the most crucial reason to deploy check constraints.

Check constraints provide better data integrity. As check constraints are always executed whenever the data in the column upon which they are defined is to be modified, the business rule is not bypassed during ad hoc processing and dynamic SQL. When business rules are enforced using application programming logic instead, the rules cannot be checked during ad hoc processes.

Check constraints provide better data integrity.

Check constraints promote consistency. Because they are implemented once—in the table DDL—each constraint is always enforced. Constraints written in application logic, on the other hand, must be executed by each program that modifies the data to which the constraint applies. This can cause code duplication and inconsistent maintenance, resulting in inaccurate business rule support.

Typically, check constraints coded in DDL will outperform the corresponding application code to perform the same edit checking. The overall impact of check constraints is to increase application development productivity while at the same time improving data integrity.

Check Constraint Examples

Check constraints enable the DBA or database designer to specify more robust data integrity rules directly into the database. Consider the following example:

[Click here to view code image](#)

```
CREATE TABLE EMP
(empno          INTEGER          PRIMARY KEY,
 CONSTRAINT check_empno
 CHECK (empno BETWEEN 100 and 25000),
emp_address    VARCHAR(70),
emp_type       CHAR(8)
 CHECK (emp_type IN ('temp', 'fulltime',
'contract'))),
emp_dept       CHAR(3)          NOT NULL WITH DEFAULT,
salary         DECIMAL(7,2)     NOT NULL
 CONSTRAINT check_salary
 CHECK (salary < 50000.00),
commission     DECIMAL(7,2),
bonus          DECIMAL(7,2)
```

```
) IN db.ts;
```

This CREATE statement for the EMP table contains three different check constraints:

- The name of the first check constraint is check_empno. It is defined on the empno column. The constraint ensures that the empno column can contain values that range from 100 to 25,000 (instead of the domain of all valid integers).
- The second check constraint for this table is on the emp_type column. This is an example of an unnamed constraint and, though possible, is not recommended. It is best to always provide an explicit constraint name in order to ease identification and administration. This specific constraint restricts the values that can be placed into emp_type to “temp,” “fulltime,” and “contract”; no other values would be accepted.
- The last check constraint on this table is named check_salary. It effectively ensures that no employee can have a salary of more than \$50,000. (Who wants to work there?)

This first check constraint example depicts column-level constraints. However, check constraints also may be coded at the table level. A column-level check constraint is defined in the DDL immediately after the column. Appropriately enough, a table-level check constraint is defined after all of the columns of the table have already been defined.

It is quite common for business data integrity rules to require access to multiple columns within a single table. When this situation occurs, it is wise to code the business rule into a check constraint at the table level instead of at the column level. Of course, any column-level check constraint can also be defined at the table level as well. In terms of functionality, there is no difference between an integrity constraint defined at the table level and the same constraint defined at the column level. Let’s augment our sample table DDL to add several table-level check constraints:

[Click here to view code image](#)

```
CREATE TABLE EMP
(empno          INTEGER          PRIMARY KEY,
 CONSTRAINT check_empno
 CHECK (empno BETWEEN 100 AND 25000),
 emp_address    VARCHAR(70),
 emp_type       CHAR(8))
```

```

        CHECK (emp_type IN ('temp', 'fulltime',
'contract')),
emp_dept      CHAR(3)          NOT NULL WITH DEFAULT,
salary        DECIMAL(7,2)    NOT NULL
        CONSTRAINT check_salary
        CHECK (salary < 50000.00),
commission    DECIMAL(7,2),
bonus         DECIMAL(7,2)
        CONSTRAINT comm_vs_salary
        CHECK (salary > commission),
        CONSTRAINT comm_bonus
        CHECK (commission=0 OR bonus=0),
) IN db.ts;

```

Now the CREATE statement for the EMP table has been modified to contain two table-level check constraints having the following ramifications:

- The first table-level check constraint for the EMP table is named `comm_vs_salary`. This constraint ensures that no employee can earn more in commission than in salary.
- The second table-level check constraint is named `comm_bonus`. This constraint will ensure that an employee cannot earn both a commission and a bonus.

Nulls and Other Potential Problems

An additional consideration for check constraints is the relational null. Any nullable column that is defined with a check constraint can be set to null. When the column is set to null, the check constraint evaluates to unknown. Because null indicates the lack of a value, the presence of a null will not violate the check constraint.

Any nullable column that is defined with a check constraint can be set to null.

Running database utilities can cause problems with check constraints. For example, depending on the DBMS, the LOAD utility may or may not enforce check constraints as data is being loaded to the table. If constraints are not enforced, data can be loaded that does not conform to the check constraints defined on the table, causing data integrity problems. If constraints are enforced during the LOAD process, someone will have to manually edit the rows that were not loaded so that they adhere to the check constraint

definitions on the table.

Another potential problem with check constraints is inconsistent coding from table to table. It is likely that similar columns of the same data type and length will exist in multiple tables throughout the database. If all of these columns must adhere to the same edit-check requirements, the DBA must create the same check constraint on each column across multiple tables. It is possible that errors can be made. Refer to the sidebar “[Rules](#)” for a discussion of a different type of constraint available in Sybase and Microsoft SQL Server that helps to alleviate this problem.

Rules

Both Sybase and Microsoft SQL Server offer a special type of check constraint called a *rule*. Although rules are similar to check constraints, they are “freestanding” database objects.

Like a check constraint, a rule defines the parameters for data validation. Whenever data is inserted or updated, the rule is checked to ensure that the data modification complies with the rule. Both columns and UDTs can have rules defined for them.

After a rule has been created, it must be bound to columns and UDTs before it has any effect on the database. Binding is accomplished using the SP_BINDRULE system procedure. For example, the following statements can be used to create a rule that forces the state code to represent only Illinois, Wisconsin, or Indiana, and then binds the rule to the state column in two different tables:

[Click here to view code image](#)

```
CREATE rule state_rule AS @state IN ("IL", "WI", "IN")
EXEC sp_bindrule "state_rule", "publishers.state"
EXEC sp_bindrule "state_rule", "authors.state"
```

Once a rule is bound to a column or user-defined data type, it will function like a check constraint. Whenever data values are inserted or updated, the rule will be invoked to ensure that the values conform to the rule definition. Rules can be removed from columns or user-defined data types by using the SP_UNBIND system procedure when the rule is no longer required.

Rules can be more reliable than check constraints because a rule need

be created only once, after which it is bound to the appropriate columns and data types. So, if one constraint applies to multiple columns, less work is required with a rule than with a check constraint. The check constraint must be explicitly coded within the DDL of every table that must be checked.

Microsoft has indicated that rules will be removed in a future version of Microsoft SQL Server. For this reason it is probably a good idea to avoid creating any new rules and focus on using constraints instead. Sybase has not indicated that rules will be deprecated from ASE, though. Additionally, it is possible to define unworkable check constraints on a table. Relational DBMS products provide the additional capability to specify column defaults. When a row is inserted or loaded into a table and no value is specified for the column, the column will be set to the value that has been identified in the column default specification. For example, the DBA could define a default for the emp_type column of our sample EMP table as follows:

[Click here to view code image](#)

```
emp_type CHAR(8)    DEFAULT 'fulltime'  
        CHECK (emp_type IN ('temp', 'fulltime',  
        'contract')), ...
```

If a row is inserted without specifying an emp_type, the column will default to the value “fulltime.” However, most DBMSs are not designed to perform semantic checking on constraints and defaults. For example, the DBMS would allow the DBA to define defaults that contradict check constraints. Furthermore, it is possible for the DBA to define check constraints that contradict one another. Care must be taken to avoid creating this type of problem. Examples of contradictory constraints follow.

In the following case, no value is both greater than 10 and less than 9, so nothing could ever be inserted into the table.

```
CHECK (empno > 10 AND empno <9)
```

In the next case, the default value is not one of the permitted emp_type values according to the defined constraint. No defaults would ever be inserted, so a value must always be provided when rows are inserted.

[Click here to view code image](#)


```
emp_type CHAR(8) DEFAULT 'new'  
        CHECK (emp_type IN ('temp', 'fulltime',  
        'contract')), ...
```

In this final case, the constraints are redundant. No logical harm is done, but both constraints will be checked, thereby impacting the performance of applications that modify the table in which the constraints exist.

```
CHECK (empno > 10)  
CHECK (empno >= 11)
```

Of course, other potential semantic problems could occur, such as

- When the parent table indicates a neutralizing DELETE RI constraint but a check constraint is defined on the child table stating CHECK (COL1 IS NOT NULL)
- When two constraints are defined on the same column with contradictory conditions
- When the constraint requires that the column be null, but the column is defined as not null

Check constraints provide a very powerful vehicle for supporting business rules in the database. They cannot be bypassed and therefore provide better data integrity than corresponding logic programmed into the application. It is a wise course of action for DBAs to define check constraints in the database instead of coding similar edit-checking logic into application programs.

Triggers

[*Triggers*](#) are event-driven specialized procedures that are attached to database tables. A trigger is a piece of code that is executed in response to a data modification statement, that is, an INSERT, UPDATE, or DELETE. To be a bit more precise, triggers are event-driven specialized procedures that are stored in, and managed by, the DBMS. Each trigger is attached to a single, specified table. Triggers can be thought of as an advanced form of rule or constraint that is written using an extended form of SQL. A trigger cannot be directly called or executed; it is automatically executed (or “fired”) by the RDBMS as the result of an action—a data modification to the associated table.

A trigger is a piece of code that is executed in response to a data modification statement.

Once a trigger is created, it is always executed when its “firing” event occurs (UPDATE, INSERT, or DELETE). Therefore, triggers are automatic, implicit, and non-bypassable. A database with triggers defined on it is sometimes called an *active database* because simple data modification results in additional actions taking place—as defined in the trigger code.

Triggers are similar to stored procedures. Both consist of procedural logic that is stored in or managed by the database. However, stored procedures are not event driven and are not attached to a specific table. A stored procedure is explicitly executed by invoking a CALL to the procedure (instead of being implicitly executed like triggers). Additionally, a stored procedure can access many tables without being specifically associated to any of them.

Triggers are useful for implementing code that must be executed on a regular basis due to a predefined event. Using triggers, you can eliminate scheduling and data integrity problems because the trigger will be fired whenever the triggering event occurs. You need not remember to schedule or code an activity to perform the logic in the trigger—it happens automatically.

Triggers can be implemented for many practical uses. Quite often, it is impossible to code business rules into the database by using only DDL. For example, most DBMSs do not support complex constraints (only value-based CHECK constraints) or certain types of referential constraints (such as pendant DELETE processing or cascading UPDATES). Using triggers, a very flexible environment is established for implementing business rules and constraints in the DBMS. This is important because having the business rules in the database ensures that everyone uses the same logic to accomplish the same process.

Triggers are quite flexible and can be coded to suit many purposes. For example, triggers can

- Access and modify other tables
- Print informational messages
- Specify complex restrictions

For example, consider a basic student enrollment database such as might be used by a college or university. A course can be attended by many students, and a student can attend many courses. Triggers can be used to support scenarios such as the following:

- Perhaps the university places a limit, say, five, on the number of

courses in which a student can enroll per semester. A trigger can be coded to check that no more than five rows are inserted into the course enrollment table for the same student. Similarly, if a course size limit is required, a trigger can be coded to check that no more students than the limit are inserted into the course enrollment table.

- A trigger can be created to ensure that each professor is assigned to at least one course each semester, or to limit the number of courses each professor can teach each semester.
- Triggers can be used to perform calculations such as ensuring that the student has paid the proper tuition for the semester.
- Triggers can be coded to ensure that the student has successfully completed prerequisite courses.

The number of business rules that can be implemented using triggers is truly limited only by your imagination, or more appropriately, your business needs.

The number of business rules that can be implemented using triggers is limited only by your imagination.

When Does a Trigger Fire?

Triggers can be coded to fire at two different times: before the firing activity occurs or after the firing activity occurs. A “before” trigger executes before the firing activity occurs; an “after” trigger executes after the firing activity occurs.

Some DBMSs support both “before” and “after” triggers, whereas others support only one type of trigger. Knowing how the triggers in your database function is imperative. Without this knowledge, properly functioning triggers cannot be coded, supported, or maintained effectively.

Consider, for example, the case where the firing activity occurs before the trigger is fired. In other words, the UPDATE, INSERT, or DELETE occurs first, and as a result of this action, the trigger logic is executed. If necessary, the trigger code may be able to roll back the data modification. What if the trigger is fired before the actual firing event occurs? In this situation, a rollback would not be required for the firing event code because it did not yet occur. However, a rollback may be required for any data modifications that occurred prior to this firing event within the same transaction.

If multiple triggers are coded on the same table, which trigger is fired first? It can make a difference as to how the triggers are coded, tested, and maintained. The rule for order of execution depends on the DBMS. For example, DB2 fires triggers of the same type on the same table in the order in which they were created. If two DELETE triggers exist on the same table, the one that was physically created first would be executed first. Be sure to understand how your DBMS fires triggers and keep the rules in mind as you make changes to your database. If you need to drop a table and recreate it to implement a schema change, make sure you create the triggers in the desired (same) order to keep the functionality the same.

The rule for order of trigger execution depends on the DBMS.

Nested Triggers

As we've already learned, a trigger is fired by an INSERT, UPDATE, or DELETE. However, a trigger also can contain INSERT, UPDATE, and DELETE statements within itself. Therefore, a data modification fires a trigger that can cause another data modification that fires yet another trigger. When a trigger contains INSERT, UPDATE, and/or DELETE logic, the trigger is said to be a *nested trigger*.

Most DBMSs, however, place a limit on the number of nested triggers that can be executed within a single firing event. If this were not done, it could be quite possible for triggers to fire ad infinitum until all of the data is removed from an entire database.

If RI is combined with triggers, additional cascading UPDATES and/or DELETES can occur. If an UPDATE or DELETE results in a series of additional UPDATES or DELETES that need to be propagated to other tables, the UPDATE or DELETE triggers for the second table also will be activated. This combination of multiple triggers and referential constraints is capable of setting a cascading effect into motion, which can result in multiple data changes. DB2 limits this cascading effect to 16 levels in order to prevent endless looping; if more than 16 levels of nesting occur, the transaction is aborted. Other DBMSs provide options to enable and disable cascading triggers.

The ability to nest triggers provides an efficient method for implementing automatic data integrity. Because triggers generally cannot be bypassed, they provide an elegant solution to the enforced application of business rules. Use

caution, however, to ensure that the maximum trigger nesting level is not reached. Failure to heed this advice can create an environment where certain types of updates cannot occur at all.

Nesting triggers provides an efficient method for implementing automatic data integrity.

Using Triggers to Implement Referential Integrity

One common use for triggers is to support RI. Triggers can be coded, in lieu of declarative RI, to support any type of referential constraint you wish to place on the database. Of course, when you use triggers, it necessitates writing procedural code for each rule for each constraint, whereas declarative RI constraints are coded in the DDL that is used to create relational tables.

Triggers can be coded to support any type of referential constraint.

Additional guidance on using triggers to support RI can be found later in this chapter.

Transition Variables and Tables

In order to use triggers to support RI rules, it is sometimes necessary to know the values impacted by the action that fired the trigger. For example, consider the case where a trigger is fired because a row was deleted. The row, and all of its values, has already been deleted because the trigger is executed after its firing action occurs. But if this is the case, how can we ascertain if referentially connected rows exist with those values? We may need to access it in its original, unmodified format.

Each DBMS has a different method of accessing versions of the data both before and after the modification from inside a trigger. In Sybase and Microsoft SQL Server, triggers use tables called `INSERTED` and `DELETED`; in DB2 the triggers specify `OLD` and `NEW`. Regardless of the specific nomenclature, these tables are generically referred to as *transition tables* (see [Figure 13.2](#)).

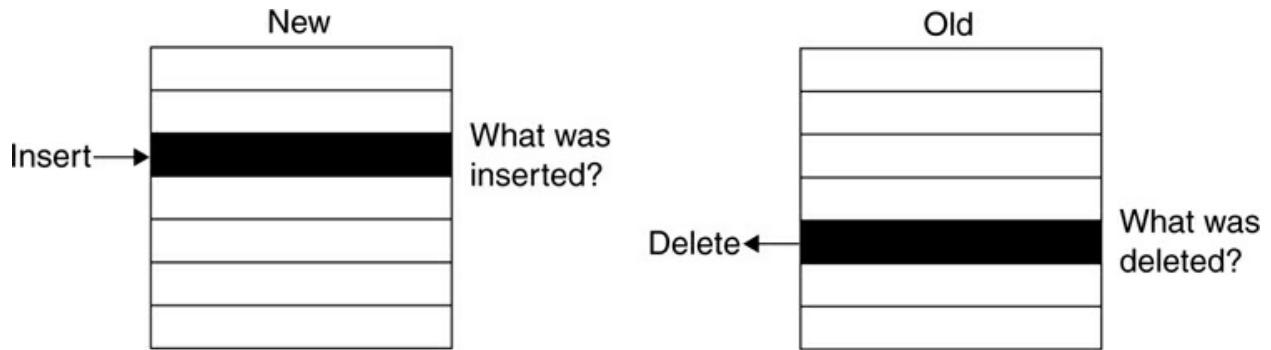


Figure 13.2. Transition tables

Each trigger has two tables, one with a “before” image of the data and another with an “after” image of the data. Transition tables are accessible only from triggers. For Sybase and Microsoft SQL Server, the transition tables operate as follows:

Transition tables are accessible only from triggers.

- When an INSERT occurs, the inserted table contains the rows that were just inserted into the table to which the trigger is attached.
- When a DELETE occurs, the deleted table contains the rows that were just deleted from the table to which the trigger is attached.
- When an UPDATE occurs, it is treated as a DELETE followed by an INSERT, so the inserted table contains the new values for the rows that were just updated in the table; the deleted table contains the old values (before they were updated) for the updated rows.

The DB2 and Oracle implementation of transition tables (or variables) is simpler and works as follows:

- When an INSERT occurs, the values inserted can be referenced with the NEW transition variable.
- When a DELETE occurs, the values deleted can be referenced with the OLD transition variable.
- When an UPDATE occurs, the inserted values are referenced using NEW, and the deleted values are referenced using OLD.

A Sample Trigger

Sometimes examining a sample piece of code can help to clarify concepts. So, let’s take a quick look at a (DB2) trigger:

[Click here to view code image](#)

```
CREATE TRIGGER salary_update
  BEFORE UPDATE OF salary
  ON emp
  FOR EACH ROW MODE DB2SQL
  WHEN (new.salary > (old.salary * 1.5))
  BEGIN atomic
    SIGNAL SQLSTATE '75001' ('Raise exceeds 50%');
  END;
```

This trigger ensures that the value of the salary column cannot be increased by more than 50 percent for any individual update. Note the use of transition variables to check the newly updated salary against the previous, old value for salary.

Trigger Granularity

Because SQL is a set-level language, any single SQL statement can impact multiple rows of data. For example, one DELETE statement can actually cause zero, one, or many rows to be removed. Triggers need to take this into account.

Therefore, there are two levels of granularity that a trigger can have: statement level or row level. A *statement-level trigger* is executed once upon firing, regardless of the actual number of rows inserted, deleted, or updated. A *row-level trigger*, once fired, is executed once for each and every row that is inserted, deleted, or updated.

A trigger can have statement-level or row-level granularity.

Different business requirements will drive the type of trigger granularity that should be chosen. But keep in mind that a row-level trigger can have a significant impact on performance because it may have to run multiple times for every SQL data modification statement issued.

INSTEAD OF Triggers

An additional type of trigger, an INSTEAD OF trigger, is significantly different from the three other types of triggers. INSTEAD OF triggers are defined on views and contain the logic required to enable nonupdateable views to become updateable. As such, there are INSERT, UPDATE, and DELETE types of INSTEAD OF triggers.

INSTEAD OF triggers can be created to make nonupdateable views updateable.

Typically, a view that consists of multiple base tables cannot be updated; however, with an INSTEAD OF trigger, this problem can be overcome. INSTEAD OF triggers can be created to direct inserts, updates, and deletes to the appropriate underlying tables of the view. When INSTEAD OF triggers are used, your application code does not have to include complex algorithms to specify which operations should be performed against views and which should be performed against base tables. Instead, all actions are performed against the view, and the activated trigger determines which underlying base tables are to be impacted.

One way to think of INSTEAD OF triggers is that they contain the inverse of the logic in the body of the view. If the view joins tables, the trigger should break the join apart to modify the correct data. Only one INSTEAD OF trigger is allowed for each type of operation on a given subject view: that is, one for inserts, one for updates, and one for deletes. Therefore, you can have a grand total of three INSTEAD OF triggers per view.

Trigger Synopsis

Triggers are a powerful feature of modern DBMS products. Triggers enable the DBA to create non-bypassable, event-driven logic that is intrinsically intermingled with data. Before implementing triggers, though, be sure to be aware of the additional DBA duties required to support an active database environment. These issues were discussed in the “[Procedural DBAs: Managing Database Logic](#)” section of [Chapter 1](#).

Referential Integrity

Referential integrity is a method for ensuring the “correctness” of data within a DBMS. People tend to oversimplify RI, stating that it is merely the identification of relationships between relational tables. It is actually much more than this. However, the identification of the primary and foreign keys that constitute a relationship between tables *is* a component of defining referential integrity.

Referential integrity is a method for ensuring the “correctness” of data.

RI embodies the *integrity* and *usability* of a relationship by establishing rules that govern that relationship. The combination of the primary and foreign key columns and the rules that dictate the data that can be housed in those key columns is the beginning of understanding and using RI to ensure correct and useful relational databases.

RI rules, applied to each relationship, determine the status of foreign key columns when inserted or updated, and of dependent rows when a primary key row is deleted or updated. In general, a foreign key must always either contain a value within the domain of foreign key values (values currently in the primary key column) or be set to null.

The concept of RI is summarized in the following “quick and dirty” definition: RI guarantees that an acceptable value is *always* in the foreign key column. “Acceptable” is defined in terms of an appropriate value as housed in the corresponding primary key, or a null.

Referential integrity guarantees that an acceptable value is always in the foreign key column.

The combination of the relationship and the rules attached to that relationship is referred to as a [referential constraint](#). The rules that accompany the RI definition are just as important as the relationship.

Two other important RI terms are *parent table* and *child table*. For any given referential constraint, the *parent table* is the table that contains the primary key, and the *child table* is the table that contains the foreign key. Examine [Figure 13.3](#). The parent table in the employed-by relationship is the DEPT table. The child table is the EMP table. So the primary key (say, DEPT-NO) resides in the DEPT table and a corresponding foreign key of the same data type and length, but not necessarily the with same column name, exists in the EMP table.




Figure 13.3. Parent and child tables

Three types of rules can be attached to each referential constraint: an INSERT rule, an UPDATE rule, and a DELETE rule. Let's see how these rules govern a referential constraint.

Three types of rules can be attached to each referential constraint.

INSERT Rule

The *INSERT rule* indicates what will happen if you attempt to insert a value into a foreign key column without a corresponding primary key value in the parent table. There are two aspects to the RI INSERT rule:

1. It is *never* permissible to insert a row into a dependent table with a foreign key value that does not correspond to a primary key value. This is known as the *restrict-INSERT rule*.
2. Whether actual values *must* be specified instead of nulls.

For each relationship, the DBA must decide whether the foreign key value(s) must be specified when the row is initially inserted into the table. To determine this, ask the following question: "Does it make sense, in business terms, to know the primary key value in the parent table when adding a dependent row?"

If a foreign key value is specified, it must be equal to one of the values currently in the primary key column of the parent table. This implements the restrict-INSERT rule. If a foreign key value is optional, it can be set to null.

UPDATE Rule

The basic purpose of the *UPDATE rule* is to control updates such that a foreign key value cannot be updated to a value that does not correspond to a primary key value in the parent table. There are, however, two ways to view the update rule: from the perspective of the foreign key and from that of the primary key.

Foreign key perspective. Once you have assigned a foreign key to a row, either at insertion or afterward, you must decide whether that value can be changed. Again, this is determined by looking at the business definition of the relationship and the tables it connects. However, if you permit a foreign key value to be updated, the new value must either be equal to a primary key value currently in the parent table or be null.

Primary key perspective. If a primary key value is updated, three options exist for handling foreign key values:

- *Restricted UPDATE.* The modification of the primary key column(s) is not allowed if foreign key values exist.
- *Neutralizing UPDATE.* All foreign key values equal to the primary key value(s) being modified are set to null. Of course, neutralizing UPDATE requires that nulls be permitted on the foreign key column(s).
- *Cascading UPDATE.* All foreign key columns with a value equal to the primary key value(s) being modified are modified as well.

DELETE Rule

Referential integrity rules for deletion define what happens when an attempt is made to delete a row from the parent table. Similar to the primary key perspective of the UPDATE rule, three options exist when deleting a row from a parent table:

- *Restricted DELETE.* The deletion of the primary key row is not allowed if a foreign key value exists.
- *Neutralizing DELETE.* All foreign key values equal to the primary key value of the row being deleted are set to null.
- *Cascading DELETE.* All foreign key rows with a value equal to the primary key of the row about to be deleted are deleted as well.

Pendant DELETE

The final type of referential constraint is the *pendant DELETE*. This special type of referential integrity deals with the treatment of parent table rows when no foreign keys from the child table refer back to the primary key. Pendant DELETE RI specifies that the parent table row be deleted after the last foreign key row that refers to it is deleted.

Pendant DELETE processing cannot be implemented using declarative RI. However, triggers can be used to code the program logic to check for this condition and execute the deletion of the primary key row.

Pendant DELETE processing cannot be implemented using declarative RI.

Please see [Table 13.1](#) for a summary of referential integrity and the rules that apply to its application.

Table 13.1. Referential Integrity Rules

RI Rule	Description
DELETE RESTRICT	If any rows exist in the dependent table, the primary key row in the parent table cannot be deleted.
DELETE CASCADE	If any rows exist in the dependent table, when the primary key row in the parent table is deleted, all dependent rows are also deleted.
DELETE NEUTRALIZE	If any rows exist in the dependent table, when the primary key row in the parent table is deleted, the foreign key column(s) for all dependent rows are set to NULL as well.
UPDATE RESTRICT	If any rows exist in the dependent table, the primary key column(s) in the parent table cannot be updated.
UPDATE CASCADE	If any rows exist in the dependent table, the primary key column(s) in the parent table are updated, and all foreign key values in the dependent rows are updated to the same value.
UPDATE NEUTRALIZE	If any rows exist in the dependent table, the primary key row in the parent table is deleted, and all foreign key values in the dependent rows are updated to NULL as well.
INSERT RESTRICT	A foreign key value cannot be inserted into the dependent table unless a primary key value already exists in the parent table.
FOREIGN KEY UPDATE RESTRICTION	A foreign key cannot be updated to a value that does not already exist as a primary key value in the parent table.
PENDANT DELETE	When the last foreign key value in the dependent table is deleted, the primary key row in the parent table is also deleted.

Setting Up Relationships

A declarative referential constraint is added by coding the primary key in the parent table and one or more foreign keys in dependent tables. Constraints can be added using the CREATE TABLE and ALTER TABLE statements. When implementing declarative referential integrity between a parent and a dependent table, certain rules must be followed.

For the parent table:

- A primary key must be identified in the CREATE TABLE or ALTER TABLE statement.
- SQL Server automatically defines a unique clustered index for the primary key when the primary key constraint is specified in the CREATE TABLE or ALTER TABLE statement. Other DBMSs—for example, DB2—require the DBA to manually create the unique index to support the primary key.

For the dependent table:

- A foreign key that references the parent table must be identified in the CREATE TABLE or ALTER TABLE statement.
- Although not usually required, it is strongly recommended that an index be defined for the foreign key also. An index on the foreign key will not be a unique index unless you are defining a one-to-one relationship. Creating an index on the foreign key will enhance the performance of RI operations.

All constraints will be named, whether explicitly or implicitly. It is better to explicitly name each referential constraint in the DDL. Failure to do so will cause the DBMS to assign a system-generated name, making future referencing of the constraint more difficult.

Declarative RI Implementation Concerns

Sometimes a referential constraint needs to be set up within a single table. For example, a table of departments may need to record which department is subordinate to which other department. A MANAGED_BY_DEPT column may need to be a foreign key of the DEPT_NO primary key—all within the same table. A table is able to reference itself in a declarative RI constraint. This is referred to as a *self-referencing constraint*.

Constraints are checked before triggers are fired. If declarative RI constraints and triggers coexist in the same tables, be sure that they are compatible with one another. For example, you should not code a DELETE trigger to delete foreign keys in conjunction with a declarative constraint, because the declarative constraint will prohibit the trigger from ever firing.

Tables can be altered to add or drop declarative RI. Indeed, sometimes it is not possible to specify all referential constraints within the initial CREATE TABLE statements. A primary key in the parent table must exist before a foreign key can be defined that references the parent table. With self-referencing constraints, the primary key definition can be included when the table is initially created, but the foreign key will have to be added later by using an ALTER TABLE statement.

Tables can be altered to add or drop declarative RI.

DBMS Support for Referential Integrity

Keep in mind that each DBMS supports different levels of declarative RI and different options for its use. The DBA must learn which RI options are

provided by the DBMS in use and provide guidance to the application developers as to what can be handled by the DBMS and what must be coded into the application. Furthermore, the DBA must understand triggers and how they can be used to implement RI when declarative RI support is lacking in the DBMS.

Each DBMS supports different levels of declarative RI.

Referential Integrity Using Triggers

If the DBMS you are using does not provide the declarative RI functionality required by your applications, triggers can be coded in lieu of declarative RI. By using triggers, it is possible to support all conceivable RI rules. Of course, using triggers necessitates writing procedural code for each rule for each constraint. Complete referential integrity can be implemented using four types of triggers for each referential constraint:

1. A DELETE trigger on the parent table can be used to code
 - DELETE RESTRICT
 - DELETE CASCADE
 - DELETE NEUTRALIZE
2. An UPDATE trigger on the parent table can be used to code
 - UPDATE RESTRICT
 - UPDATE CASCADE
 - UPDATE NEUTRALIZE
3. An INSERT trigger on the dependent table can be used to code
 - INSERT RESTRICT
4. An UPDATE trigger on the dependent table can be used to code the restriction that a foreign key cannot be updated to a non-primary-key value.

Indeed, nested and recursive triggers can be coded to support a very robust implementation of referential integrity. Furthermore, triggers can be the only automatic method of implementing system-managed RI under certain circumstances:

- When deleted, inserted, and updated information needs to be explicitly referenced in order to determine the action to take. Triggers provide a

method of doing this, which we will discuss shortly.

- When an RI rule that is not supported by the DBMS is necessary. For example, Sybase declarative RI supports only restricted DELETES and UPDATES; the declarative RI provided by DB2 supports restricted DELETES and UPDATES, neutralizing DELETES and UPDATES, and cascading DELETES, but not cascading UPDATES. Each DBMS provides different degrees of declarative RI support.
- When pendant DELETE processing is required. This is sometimes referred to as “reverse” RI. Pendant DELETE processing is the situation where the parent row must be deleted when the last dependent row that references it is deleted. The only way to implement this type of constraint is with a trigger.

Nested and recursive triggers can be coded to support a very robust implementation of referential integrity.

Of course, if your RI needs can be satisfied using declarative RI, you should use declarative RI instead of triggers. Triggers must be hand coded and tested, meaning that they may contain bugs or function improperly. Declarative RI is built into the DBMS and is unlikely to fail (unless the DBMS itself has bugs). Additionally, declarative RI almost always outperforms trigger-based RI. So, as a rule of thumb, favor declarative RI over trigger-based RI.

Let’s examine some sample triggers to clarify how transition tables and variables are properly used. The following (SQL Server) trigger implements the cascading DELETE RI rule:

[Click here to view code image](#)

```
CREATE TRIGGER title_del
ON titles FOR DELETE
AS
IF @@rowcount = 0
    RETURN
DELETE titleauthor
    FROM titleauthor, deleted, title
    WHERE titles.title_id = deleted.title_id
RETURN
```

When a row in the parent table (titles) is deleted, the DELETE is cascaded

to the dependent table (titleauthor). A second example follows that implements the restricted INSERT RI rule. When a row in the dependent table (titleauthor) is inserted, we must first check to see whether a viable primary key exists in the parent table (titles).

Let's look at another example. The following trigger is coded to implement the restricted INSERT RI rule. When a row in the dependent table (titleauthor) is inserted, we must first check to see whether a viable primary key exists in the parent table (titles).

[Click here to view code image](#)

```
CREATE TRIGGER title_ins
ON titleauthor FOR INSERT
AS
DECLARE @rc int
SELECT @rc = @@rowcount
IF @rc = 0
    RETURN
IF (select count(*)
    FROM titles, inserted
    WHERE titles.title_id = inserted.title_id) != @rc
BEGIN
    RAISERROR 20001 "Invalid title: title_id
    does not exist on titles table"
    ROLLBACK transaction
    RETURN
END
RETURN
```

A final (SQL Server) example depicts neutralizing updates:

[Click here to view code image](#)

```
CREATE TRIGGER title_upd
ON titles FOR UPDATE
AS
IF UPDATE (title_id)
    IF (select count(*)
        FROM deleted, titles
        WHERE deleted.title_id = title.titleid) = 0
    BEGIN
        UPDATE titleauthor
            SET titleauthor.titleid = null
        FROM titleauthor, deleted
        WHERE titleauthor.titleid = deleted.title_id
```


END
RETURN

The first check is to see whether the `title_id` was actually updated. If so, the trigger checks to make sure that the `title_id` was not updated to the same value it previously held. If it was, the neutralizing update should not occur. If these two checks are passed, the update occurs. When a row in the parent table (`titles`) is updated, we check to see whether any corresponding rows exist in the dependent table (`titleauthor`). If so, the foreign key columns must be set to null.

User-Managed versus System-Managed RI

System-managed declarative RI has not always been available, and your installation may have applications with user-managed RI already in place. It may be necessary to support both user- and system-managed RI in this situation.

It may be necessary to support both user- and system-managed RI.

Furthermore, even though system-managed RI is now available, sometimes user-managed RI may be a more appropriate solution. One such instance is when it is *always* necessary for applications to access the parent and dependent tables (even when system-managed RI is implemented). For example, consider an application program that always inserts the order row into the `ORDR_TAB` (parent) table before inserting the order item rows into the `ORDR_ITEM_TAB` (dependent) table; another application always accesses the rows in the `ORDR_ITEM_TAB` table for historical information before deleting them and then deleting the parent row from the `ORDR_TAB` table. Since these applications already access both tables, the additional overhead of system-managed RI may not be worthwhile.

However, the benefit of system-managed RI is that the integrity of the data is enforced during ad hoc access, such as interactive SQL and data warehouse queries. When RI is maintained only in programs, data integrity violations can occur if data modification is permitted outside the scope of the application programs that control RI. It is usually a wise move to implement system-managed rather than user-managed RI. But remember, most DBMSs provide two methods of implementing system-managed RI: declarative

constraints and triggers.

RI versus Check Constraints versus Program Logic

Referential integrity, check constraints, and writing application program logic are all valid methods of enforcing data integrity. This chapter illustrated the functionality and strengths of each. But let's stop to examine a particular requirement and the impact of implementing it using each of these methods.

Consider comparing the consequences of enforcing a domain. For example, suppose there are three valid employee types when the database is created: T (temporary), F (full-time), and C (contract). Subsequently, after implementation, a fourth employee type is added, say, P (part-time). It will be easier to support the new value using RI by creating an employee type table to store the values. To support the new type, all that is required is inserting a row into the employee type table—the referential constraints will now allow the new value.

To support this new value using a check constraint, the DBA most likely will be required to drop and recreate the check constraint. During the period when the check constraint does not exist, an outage will be required, otherwise data integrity violations could occur. Certainly, such a scenario should be avoided when a 24/7 environment is required.

Finally, if program logic is required to support a new value for the domain, a programmer must be deployed to modify the program, test it, and then migrate it to the production environment. This is a lot more work than either of the previous two scenarios.

As a rule of thumb, then, use the facilities of the DBMS to enforce data integrity unless a very compelling reason exists not to do so.

Use the facilities of the DBMS to enforce data integrity wherever possible.

RI Rules of Thumb

Regardless of the type of RI you plan to implement in your databases, you should heed the following rules of thumb.

Be aware that there are certain situations in which referential integrity can be bypassed. This can cause severe data integrity problems as well as significant confusion. One common problem area for most DBMSs is the

BULK LOAD utility. Loading data into database tables without checking foreign key references speeds up the load process. However, this also means that data integrity problems may be introduced during the bulk data load process.

Additionally, some DBMSs provide SQL switches that turn RI on and off. Avoid using these. Turning off RI is very likely to cause data integrity problems. Programs and ad hoc data modifications are likely to insert or update foreign key columns to values that do not conform to any existing primary key value in the parent table.

Primary and foreign key columns can have different names, null attribute qualifiers (e.g., NULL versus NOT NULL), and default values. The column attributes, for example, CHAR(5), must be the same. It is not possible to create a declarative constraint between two columns with different attributes. Likewise, though possible, it is not wise to create a trigger-based constraint between columns with differing attributes.

Keep in mind that when multiple relationships exist for the dependent row, they must all be verified before the row can be inserted. It is possible for an insert to pass three of four referential constraints, but to fail on the fourth.

When composite keys are used for a “primary key–foreign key” relationship, a single row must exist in the parent table with key values that match all the columns of the foreign key for the row being inserted into the dependent table.

In general, avoid using declarative RI to create very large groups of referentially related tables. The larger the group of referentially related tables, the more difficult they become to manage. Keep in mind that referentially related tables must be backed up and recovered together to keep the relationships intact.

Referentially related tables must be backed up and recovered together.

Consider using check constraints instead of referential constraints for code tables and reference tables that are small and relatively static in nature. If the number of values in a table is small, a check constraint will probably outperform referential constraints. For example, consider a table that contains status codes for projects, such as the one in [Table 13.2](#). Instead, we could

simply code a check constraint consisting of those values on any child tables, for example:

[Click here to view code image](#)

```
CREATE TABLE PROJECT
(projno      INTEGER      PRIMARY KEY,
 deptno      INTEGER      NOT NULL,
 mg_empno    INTEGER,
 status      INTEGER      NOT NULL
      CHECK (status BETWEEN 0 AND 3),
 start_date  DATE,
) IN db.ts;
```

Table 13.2. PROJECT_STATUS Codes

STATUS_CODE	Description
0	Project not yet started
1	Project begun
2	Project complete
3	Project past deadline

The check constraint on the status column performs the same function as creating a foreign key on the status column to the STATUS_CODE column in the PROJECT_STATUS table (see [Table 13.2](#)). And it will be more efficient. However, if the project status codes change frequently, it will be more difficult to maintain the values using check constraints than it would be to add values to the PROJECT_STATUS table.

Temporal Database Systems

Recall from [Chapter 4](#) that a temporal DBMS makes it possible to store different database states and to query the data “as of” those different states. A DBMS that provides temporal support must be defined differently to manage the temporal integrity of the data. Data stored in a temporal database differs from traditional data in that a time period is attached to the data to indicate when it was valid or changed in the database.

Temporal support makes it possible to store different database states and to query the data “as of” those different states.

A temporal database implementation requires setting up constraints for the table indicating the valid time range for each row. This typically is done by specifying the business or system time period start and end dates. For example, consider the following DDL to create a new table with a business time specification:

[Click here to view code image](#)

```
CREATE TABLE COURSE
(COURSENO    INTEGER          NOT NULL,
 TITLE      VARCHAR(20)     NOT NULL,
 CREDITS     SMALLINT        NOT NULL WITH DEFAULT 3,
 PRICE      DECIMAL(7,2)    NOT NULL,
 BUS_START  DATE            NOT NULL,
 BUS_END    DATE            NOT NULL,

 PERIOD BUSINESS_TIME(BUS_START, BUS_END) ,

 PRIMARY KEY(COURSENO, BUSINESS_TIME WITHOUT
 OVERLAPS)
);
```

In this example, two DATE columns are specified, indicating the beginning of the time period (BUS_START) and the end (BUS_END). The PERIOD BUSINESS_TIME clause is used to specify these two columns as identifying the time period. It is a good idea to ensure the temporal integrity of the data by enforcing BUS_START values to always be before BUS_END values. The DBMS may automatically generate an implicit constraint to enforce this condition.

The WITHOUT OVERLAPS clause is specified as part of the primary key constraint for the table. This indicates that primary key values must be unique for any point within the business time. In other words, there cannot be two states, or versions, of the course that are valid at the same time. If WITHOUT OVERLAPS is not specified, overlaps can occur.

The data, including the time period dates, must be supplied by the application. All data, including past, present, and future effective-dated rows, is maintained in the table. Temporal SQL enables queries to be issued that specify temporal predicates such as AS OF, FROM/TO, and BETWEEN. For example, the following SQL selects the price of the English 101 course offering as of April 3, 2012:

[Click here to view code image](#)

```
SELECT Price
FROM   Course FOR BUSINESS_TIME AS OF '2012-04-03'
WHERE  Title = 'ENGLISH 101';
```

The temporal constraints are helpful, too, when data is to be modified in the table. Inserting data is a simple matter of supplying the appropriate values for all required columns including the business time start and end columns. If you have defined a temporal uniqueness constraint on the table, the DBMS reviews any attempted INSERT to ensure that it does not conflict with the existing time periods.

To update data in a temporal table, you can write traditional UPDATE statements. Alternatively, you also can use the new FOR PORTION OF BUSINESS_TIME clause to restrict the UPDATE to a specific business time period. The temporal syntax is powerful because an UPDATE that spans multiple time periods is managed by the DBMS via modifying the time ranges appropriately. For example, consider the Course table data for the Introduction to DBA class as follows:

CourseNo	Title	Credits	Price	Bus_Start	Bus_End
100	INTRO TO DBA	3	300.00	2011-10-01	2012-01-01
100	INTRO TO DBA	3	350.00	2012-01-01	9999-12-31

Suppose that you need to change the price for this course for the time period from November 1, 2011, to February 2, 2012. This UPDATE spans the two existing rows for the course because it begins within the time period of the first row and ends within the time period of the second. You could, of course, issue several UPDATE and INSERT statements to adjust the data to include the wanted change. But it is simpler to use the FOR PORTION OF BUSINESS_TIME clause as follows:

[Click here to view code image](#)

```
UPDATE Course
FOR PORTION OF BUSINESS_TIME FROM '2011-11-01' TO
'2012-02-02'
SET   Price = 325.00
WHERE CourseNo = 100;
```

Both rows are impacted by the UPDATE statement because the portion of business time that is updated overlaps partially with the business period of

each row. To make this change, the DBMS splits each of the two original rows into two rows. The time periods for these rows are then adjusted accordingly to match the requirement of the UPDATE and the changed column values. (In this case, Price is applied to the appropriate rows.) After this UPDATE statement is issued, the data for CourseNo 100 now consists of four rows instead of two, each with the appropriate temporal range.

You can also DELETE data from a table with business time periods. You can use either a standard DELETE, which impacts rows without regard to time period, or you can restrict the DELETE to a specific time range using the FOR PORTION OF BUSINESS_TIME clause. Similar to the UPDATE, when data is to be deleted that is not completely contained within the specified time range, the DBMS takes appropriate measures to retain the data that must be preserved.

Temporal support offers powerful constraints for managing time in the database. Keep in mind, though, that not every DBMS offers temporal support. DB2 for z/OS Version 10 is an example of one that does.

Summary

Today's DBMSs provide a wealth of features to support data integrity. Because one of the major problems plaguing production systems today is data quality, it is imperative that DBAs understand, implement, and administer DBMS features such as constraints, referential integrity, and triggers in their database designs. Failure to do so can be a prescription for disaster.

However, no DBMS can ensure the integrity of its data 100 percent reliably all of the time. Other methods—programs, manual inspections, automated data cleansing, and data quality products—may be needed to assure the quality of specific databases.

No DBMS can ensure the integrity of its data 100 percent reliably all of the time.

However, every DBA should take advantage of the mechanisms provided by the DBMS to ensure data integrity. When DBMS-provided methods are used, fewer data integrity problems are likely to be found in the database. Moreover, having fewer data integrity problems means higher-quality databases and more proficient end users.

Review

1. Define what is meant by *entity integrity*.
2. Describe the difference between database structural integrity and semantic data integrity.
3. What is the preferred method of implementing referential integrity in a relational database?
4. What is a trigger, and how can triggers be used to enhance data integrity?
5. What is the difference between a check constraint and a rule?
6. What are transition variables and tables, and what benefit do they provide?
7. Name and describe four types of database structural integrity problems that may be encountered by a DBA.
8. What is a user-defined data type, and how can it be used to enhance data integrity?
9. Describe the three RI DELETE rules and the impact of each on foreign key rows in the child table when a DELETE is issued for the primary key row in the parent table.
10. Describe the purpose of the start and end date/time columns within the context of a temporal database.

Suggested Reading

Chisholm, Malcolm. *Managing Reference Data in Enterprise Databases: Binding Corporate Data to the Wider World*. San Diego, CA: Academic Press/Morgan Kaufmann (2001). ISBN 1-55860-697-1.

Codd, E. F. *The Relational Model for Database Management, Version 2*. Reading, MA: Addison-Wesley (1990). ISBN 0-201-14192-2.

Date, C. J., et al. *Temporal Data and the Relational Model*. San Francisco, CA: Morgan Kaufmann (2003). ISBN 1-55860-855-9.

Johnston, Tom, and Randall Weiss. *Managing Time in Relational Databases*. Burlington, MA: Morgan Kaufmann (2010). ISBN 978-0-12-375041-9.

Owens, Kevin T. *Building Intelligent Databases with Oracle PL/SQL*,

Triggers, and Stored Procedures. Upper Saddle River, NJ: Prentice Hall (1996). ISBN 0-13-443631-8.

Parsaye, Kamran, et al. *Intelligent Databases*. New York, NY: John Wiley & Sons (1989). ISBN 0-471-50345-2.

Piattini, Mario, and Oscar Diaz, eds. *Advanced Database Technology and Design*. Norwood, MA: Artech House (2000). ISBN 0-89006-395-8.

Rob, Peter, and Carlos Coronel. *Database Systems: Design, Implementation, and Management*. 10th ed. Boston, MA: Course Technology (2012). ISBN 978-1-111-96960-8.

Saracco, Cynthia Maro. *Universal Database Management: A Guide to Object/Relational Technology*. San Francisco, CA: Morgan Kaufmann (1998). ISBN 1-55860-519-3.

Widom, Jennifer, and Stefano Ceri. *Active Database Systems: Triggers and Rules for Advanced Database Processing*. San Francisco, CA: Morgan Kaufmann (1996). ISBN 1-55860-304-2.

14. Database Security

Database security and protection are receiving more attention and budget from organizations due to the steady increase in data breaches and the resultant regulations (see next chapter) designed to abate them. However, database security still requires more focus and effort. According to a 2005 report from Forrester Research, 75 percent of enterprises do not have a DBMS security plan.

Data Breaches

Data breaches continue to dominate both business and IT news, with bigger and uglier data breaches being regularly announced. A recent incident came from Epsilon, an e-mail marketing firm that sends billions of e-mails annually. The breach involved the exposure of e-mails for customers of many well-known brands, including the rewards programs of Ritz-Carlton and Marriott, banks such as Citibank and Capital One, and other companies such as Kroger and Walgreens.

How bad is this problem? The Privacy Rights Clearinghouse (www.privacyrights.org/data-breach) keeps track of every data breach that is reported. According to its research, more than 544 million records were breached between January 10, 2005, and early 2012—over the course of almost 3,000 separate data breach events.

Can we put a price tag on all of that unprotected and lost data? We can try.

A Forrester Research survey of companies that had experienced a data breach concluded that the average security breach can cost a company between \$90 and \$305 per lost record. But coming up with a precise figure can be difficult because of the additional, extenuating circumstances surrounding data breaches. The cost has to factor in such details as the expenses of legal fees, call centers, lost employee productivity, regulatory fines, customer losses, stock losses, and the nebulous cost of bad publicity.

The average security breach can cost a company between \$90 and \$305 per lost record.

Another research group, Ponemon Institute, conducts an annual study of

data breaches. According to its sixth annual U.S. Cost of a Data Breach Study, data breach costs increased to \$214 per compromised record in 2010. This is a significant increase from the 2006 report, which pegged the average cost per lost customer record at \$182.

If you are still skeptical, you can always roll your own numbers using the free Web-based data loss calculator provided by Darwin Professional Underwriters, Inc., at www.tech-404.com/calculator.html.

The obvious conclusion is that data breaches are costly, even at the low end of \$90 per record. Consider a typical data breach case. On February 27, 2008, Health Net Federal Services reported that thousands of doctors in 11 states had their personal information (including Social Security numbers) openly posted on a company Web site. According to the Privacy Rights Clearinghouse, the total number of records involved was 103,000. So what did that cost? At the low end, the cost was \$9.3 million, but at the high end it ballooned to over \$31.4 million. Using the Ponemon estimate, which is sort of in the middle, the cost ranged from \$18.7 million to \$22 million.

An additional interesting data point comes from the 2011 Data Breach Investigation Report, a study conducted by the Verizon RISK Team with cooperation from the U.S. Secret Service and the Dutch High Tech Crime Unit. According to this report, 76 percent of all data breached was compromised from servers. Furthermore, database servers yielded the majority of breached data.

So the cost of a data breach can be quite steep. And database servers are favorite sources for hackers to target. Thus, it makes sense to spend some time and money up-front to better secure and protect the data in your database systems.

Database Security Basics

The basic security and authorization approach taken by DBMS vendors to secure database access is that all database resources are controlled by the DBMS. No default authorizations are given to any user just because the user logs in to the DBMS. Therefore, for a user to be able to perform any DBMS operation or function, one of the following conditions must exist:

- The user has been granted the ability to perform that function or operation, or

- That operation or function has been granted generically to all users.

All database resources are controlled by the DBMS.

Using the security features of the DBMS, the DBA can set up the environment such that only certain users or certain application programs are allowed to perform certain operations on certain data within the database. Each user's function within the organization should determine the authorized level of database access. For example, only general-ledger programmers, batch jobs, and programs can access and modify the general-ledger databases. Different checks can be established for each type of access to each type of information, and different users can be assigned different access rights to different database objects.

The operational challenge of effectively administering database security arises because setting up and managing database authorization requires technical expertise and elevated privileges. Many aspects of database security require different utilities, system procedures, and commands to implement. When users require access to multiple databases, on multiple servers distributed across different physical locations, database security administration becomes quite complicated indeed. The commands must be repeated for each database, and there is no central repository for easily modifying and deleting user security settings on multiple databases simultaneously.

Although the DBA typically is responsible for administering database security, some organizations have transferred this task to a separate security administration function that controls all of the IT security for the company. However, even in many shops where security administration is a separate entity, database security is still handled by the DBA group because database security is managed differently from a typical IT authorization scenario.

When the security administration group handles security policies, this group usually relies on third-party security software such as IBM's RACF or Computer Associates' ACF2 and Top Secret. These products automate the security function and do not require the administrator to have elevated privileges to manage security policies. However, most of these security administration products run only on mainframes. Additionally, most IT security departments are understaffed and lack the technical DBMS expertise required to administer database security. Granting untrained security

personnel the privileges required to administer database security can result in accidental disruption of database service or performance problems. So, DBAs are forced to manage database security as a component of their job.

DBAs must manage database security as a component of their job.

At a high level, database security boils down to answering four questions:

- Who is it? (authentication)
- Who can do it? (authorization)
- Who can see it? (encryption)
- Who did it? (audit)

Strong authentication is the cornerstone of any security implementation plan. It is impossible to control authorization and track usage without it. Before authorization to use database resources can be granted, a login needs to be established for each user of the DBMS. Logins are sometimes referred to as *accounts*, or user IDs. The login will have a password associated with it such that only those who know the password can use the login ID. Some DBMSs use the operating system login ID and password as the DBMS login ID and password; others require an additional login ID and password to be created specifically for database access and security.

Strong authentication is the cornerstone of any security implementation plan.

When the DBMS controls the addition of logins, the DBA is required to provide certain information about the login when it is created. Typically, other than the actual login name or ID, the following information either can or must be provided:

- *Password*—the key phrase, word, or character string associated with the new login that must be provided by the user before access to the database is permitted
- *Default database*—the name of the database to which the user will initially be connected during login
- *Default language*—the default language assigned to the login when using the DBMS if multiple languages are supported

- *Name*—the actual full name of the user associated with this login
- *Additional details*—additional details about the user for which the login has been created: e-mail, phone number, office location, business unit, and so on (this is useful for documentation purposes)

Passwords should be changed regularly over time to make it difficult for hackers and miscreants to gain access to the DBMS surreptitiously. Refer to the sidebar “[Password Guidance](#)” for tips on creating useful passwords. As a DBA, you may decide to set up automated procedures—such as an e-mail notification system—to coerce users into changing their login password every month or so. Users who do not change their login password can be disabled until they call to complain. Of course, this adds to the workload of the DBA, but it does enhance the security of the DBMS.

Passwords should be changed regularly over time.

When a DBMS user no longer requires access to the DBMS, or leaves the company, the DBA should drop that user’s login from the system as soon as possible. However, this could become a complicated task—a login cannot be dropped if the person is currently using a database, or if the user owns any database objects.

For this reason it is wise to limit the database users who can create database objects to DBAs only, especially in a production environment.

Limit the database users who can create database objects to DBAs only.

In lieu of dropping a login, the DBMS may provide an option to lock the login. Locking a login prohibits the user from accessing the DBMS, but it does not actually drop the login from the system. The login can subsequently be unlocked, thereby enabling access to the server. Such a process is very useful if you simply wish to prohibit access, say, for those users who have not recently changed their login password.

Keep in mind that logins that are dropped must be created anew if they are ever required again. For this reason, follow these rules of thumb regarding login administration:

- Lock logins that may need to be reactivated.
- Drop logins that will never need to be reactivated.

Password Guidance

As a DBA, you are responsible for securing the DBMS and granting data access to authorized users. One way of assuring proper DBMS usage is to develop and disseminate tips on creating useful passwords. A password that is useful and proper will be difficult to guess.

If passwords are too simplistic or too related to some aspect of the person using the password, unscrupulous individuals may be able to guess the password and use it to surreptitiously access the database.

The following guidelines should be followed for proper password creation:

1. Avoid passwords that are too short. Each password should be at least six characters long, more if possible.
2. Each password should consist of at least a combination of alphabetic characters and numeric characters. Using other allowable symbols makes the password harder to guess.
3. Avoid creating a password that is a complete word (in either the native language of the user or any foreign language).
4. Do not embed personal statistics in the password. Street addresses, Social Security numbers, phone numbers, and the like are easily guessed and do not belong in passwords.
5. Consider concatenating two unrelated words with a symbol or number between them. For example, “toe3star” is a viable password.

Use mnemonic devices to help you remember passwords. For example, use a sentence such as “Moondance by Van Morrison is my favorite album” to remember that your password is “mbvmimfa” (the first letters of each word in the sentence). However, do not make the sentence too obvious—for example, using “My name is Craig S. Mullins” to remember “mnicism” is not a good idea because it might be quite easily guessed.

Other common and weak password archetypes include sports teams and sports celebrities (like “ChiBulls99” or “TigerWoods”), naughty words (curse words/anatomy), and names of family members. In general, if you think it can be easily guessed, avoid using it.

As a DBA, you should work with your organization’s security

administration team to create guidelines such as these and distribute them to your database users.

Some DBMSs provide additional controls and parameters on logins and passwords. For example, some DBMSs provide profile parameters for passwords that can be used to limit the following:

- Number of failed login attempts before the account is locked
- Number of days a password is valid, the grace period for changing an expired password
- Number of days the account can remain locked when the password expires
- Reusability of passwords (number of days before a password can be reused and the maximum number of times a password can be reused)

When such controls are available, be sure to use them when doling out login accounts to better secure the DBMS environment.

Keep in mind, though, that each DBMS is different, and there may be no capability to force a user to periodically change a password. If a password is never changed, the likelihood of it being compromised increases over time. If the database is able to force a periodic password change, it is often limited in its ability to completely enforce corporate password standards that reduce the risk of the password being guessed. These standards usually include minimum length and alphanumeric requirements. Most DBMSs do not even provide a simple interface by which end users can change their own password. The problem is compounded if users have accounts on multiple databases across multiple servers.

Enforce corporate password standards.

Database Users

In addition to a login account, which is required to access the DBMS, some database systems require an additional account in order to use specific databases. In this situation, a user name is created for the login account and is attached to each database required by the user. As shown in [Figure 14.1](#), the following accounts may be required:

- A *login*, sometimes called an *account*, is used to access the DBMS, or

database server. For this reason it is sometimes also known as a *server user ID*, or SUID.

- A *user name* is sometimes referred to as a *database ID*. The user name is associated with the login account. Users are required by some DBMS implementations to be set up with a database user name in order to access each database.

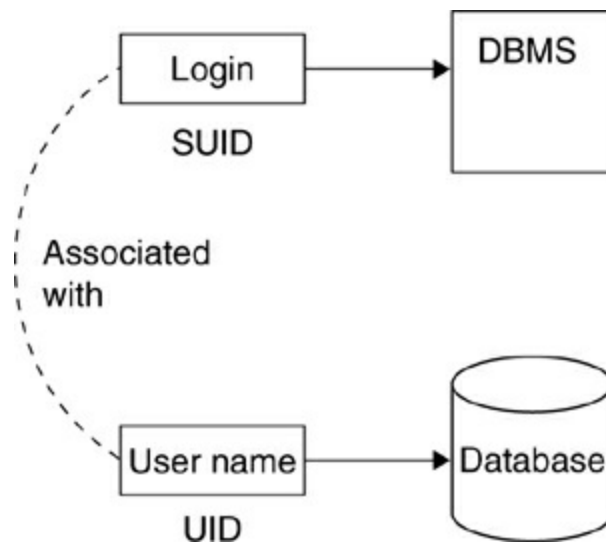


Figure 14.1. DBMS and database logins

Guest usage of a database is permitted by configuring a special user name that permits users to access the database as a GUEST. Adding a GUEST user for a database allows any login with the special user name to access the database.

Granting and Revoking Authority

The DBA controls database security and authorization using Data Control Language, or DCL. DCL is one of three subtypes of SQL. (The other two are DDL and DML.) DCL statements are used to control which users have access to which objects and commands. These statements are the manner in which database security is enacted. DCL statements comprise two basic types:

- *GRANT* assigns a permission to a database user.
- *REVOKE* removes a permission from a database user.

DCL statements comprise two basic types: GRANT and REVOKE.

The GRANT statement is issued with two accompanying lists: a list of privileges to be assigned to a list of users. To use the GRANT statement, the user must be the owner of the database object, have been granted high-level group authority, or have been given the WITH GRANT OPTION when the privilege was granted.

The WITH GRANT OPTION allows a user to pass the authority to grant privileges along to others. Generally, the use of this clause depends on whether an installation practices centralized or decentralized administration of privileges.

- *Decentralized administration* is generally easier to establish, but more difficult to control. As more and more users obtain the authority to grant privileges, the scope of authority is widened and becomes unwieldy.
- *Centralized administration* is generally easier to administer but places a burden on the centralized administrator as the sole arbiter of privileges within the environment.

Avoid issuing GRANT and REVOKE statements from within an application program. Ideally, an individual who understands the security needs of the organization—usually the DBA—grants database authorization. Furthermore, application programs designed to grant database privileges must be executed by a user who has the appropriate authority to issue the GRANTs and REVOKEs coded into the application program. This could create a loophole in your database security infrastructure.

Some DBMSs provide additional functionality for conferring or prohibiting database authority. For example, Microsoft SQL Server includes the DENY statement to specifically deny privileges to a specific user.

Types of Privileges

There are different types of privileges that can be granted and revoked from database users. Every DBMS provides certain basic types of privileges, such as the ability to access data, to create database objects, and to perform system functions. Each DBMS will also have additional types of privileges, depending on the features it supports.

The following types of privileges are commonly provided by modern DBMSs:

- [*Table*](#): to control who can access and modify the data within tables
- *Database object*: to control who can create new database objects and drop existing database objects
- *System*: to control who can perform certain types of system-wide activities
- *Program*: to control who can create, modify, and use database programs
- [*Stored procedure*](#): to control who can execute specific functions and stored procedures

Granting Table Privileges

Table privileges are granted to enable users to access tables, views, and columns within tables and views. The following privileges can be granted for tables and views:

- **SELECT**: to enable the user to select from this table/view
- **INSERT**: to enable the user to insert rows into this table/view
- **UPDATE**: to enable the user to update this table/view
- **DELETE**: to enable the user to delete rows from this table/view
- **ALL**: to enable the user to select, insert, update, and delete using this table/view

For example, to enable user7 to delete rows from the Titles table, the following statement can be issued:

```
GRANT DELETE on Titles to user7;
```

Some table privileges can be specified at the column level. Doing so can be desirable when certain users must be permitted to modify specific columns of a table but not other columns. The **SELECT** and **UPDATE** privileges can be granted for specific columns. For example, to enable user7 to update only the `au_id` column in the Titles table, the following statement can be issued:

```
GRANT UPDATE on Titles (au_id) to user7;
```

Typically, the DBA will grant table privileges to programmers in a test environment for development purposes. Additionally, programmers and end users may require table privileges on production tables for certain tasks. However, most production access should be controlled using program and

stored procedure privileges, instead of direct table privileges.

Granting Database Object Privileges

Database object privileges control which users have permission to create database structures. The actual privileges that can be granted will depend on the DBMS and the types of database objects supported. Generally, the DBMS will provide options to grant CREATE privileges on each type of database object, including databases, tablespaces, tables, indexes, triggers, defaults, and user-defined data types.

Database object privileges control which users have permission to create database structures.

For example, to enable user5 and user9 to create tables and indexes, the following statement can be issued:

```
GRANT CREATE table,  
        CREATE index  
TO user5,  
   user9;
```

The ability to create database objects is usually reserved for DBAs. If these privileges are granted to others, the number of existing database objects can be quite difficult to control. Furthermore, it becomes very difficult to track which database objects are really being used, and which were created and then abandoned. For these reasons, DBAs should keep this authority to themselves, with only rare exceptions perhaps for SAs or very skilled developers.

The ability to create database objects is usually reserved for DBAs.

Granting System Privileges

System privileges control which users can use certain DBMS features and execute certain DBMS commands. The system privileges available will vary from DBMS to DBMS but may include the ability to archive database logs, shut down and restart the database server, start traces for monitoring, manage storage, and manage database caches.

System privileges cannot be granted at the database level. System privileges are granted at a system-wide level across the DBMS. For example,

to enable user6 to start performance traces, the following statement can be issued:

```
GRANT TRACE
TO    user6;
```

System privileges cannot be granted at the database level.

System privileges should be granted with care and should usually be reserved for the DBA and SA.

Granting Program and Procedure Privileges

Granting the EXECUTE privilege gives the user permission to execute a program or a stored procedure. For example, to enable user1 and user9 to execute the stored procedure named proc1, the following statement can be issued:

```
GRANT EXECUTE on proc1
TO    user1, user9;
```

Granting privileges to users on programs and procedures is easier to control than granting privileges on individual tables and columns. The procedural logic in the program and procedure controls which specific tables and columns can be modified. Furthermore, the DBA can better maintain the integrity of production data if the only way it can be changed is programmatically.

Granting to PUBLIC

As an alternative to granting access to a database user, the DBA can choose to grant a particular authorization to PUBLIC. When authorization is granted to PUBLIC, the DBMS will allow anyone who can log in to the DBMS that particular authority. Grants made to PUBLIC cannot be given with the WITH GRANT OPTION, as everyone is in PUBLIC.

For example, to grant everyone the authority to delete rows from the titles table, the following statement can be issued:

```
GRANT DELETE on titles to PUBLIC;
```

Administering security can be a complex duty. Using the PUBLIC authority to permit blanket access to certain database objects and resources often appears easier than specifically controlling access by user. However,

DBAs should exercise caution when granting any privileges to PUBLIC.

Exercise caution when granting any privileges to PUBLIC.

Whenever a privilege is granted to PUBLIC, the DBA loses control over that database object or resource—anyone can access or use the object or resource as specified by the GRANT statement. Inevitably, users will abuse PUBLIC resources, and gaining control over the resource again will be quite difficult.

Reserve usage of PUBLIC authority for those few database objects and resources that should be available to everyone. Alternatively, PUBLIC authority can be a useful shortcut if another security mechanism is in place. For example, you could grant program authority to PUBLIC for transactions that are run under a transaction processor and use the security facilities of the transaction processor to control access.

PUBLIC Authority and the System Catalog

A common dilemma faced by DBAs is whether or not to grant PUBLIC access to the system catalog tables. On the one hand, the metadata contained in the system catalog is useful information that will be needed by most DBAs, developers, and analysts as they create and maintain applications that access production databases. On the other hand, the system catalog frequently contains sensitive information that can be exploited by hackers. For example, system catalog tables that contain authorization details can be queried to learn who has access to which tables. These tables typically contain user ID and database login information, too. Furthermore, being able to query the system catalog exposes the names of all the tables that have been created in the system, making it easier for a hacker to attempt to access data (or to attempt a SQL injection attack).

The bottom line is that you should refrain from granting blanket PUBLIC access to the system catalog tables.

Refrain from granting blanket PUBLIC access to the system catalog tables.

Revoking Privileges

The REVOKE statement is used to remove privileges that were previously granted. The syntax for REVOKE is the “flip side” of the GRANT syntax.

Additionally, privileges will be automatically revoked by the DBMS when a database object is dropped.

For example, to revoke the ability to update the `au_id` column of the `titles` table from `user7`:

```
REVOKE UPDATE on titles (au_id) from user7;
```

Revoking a `PUBLIC` privilege will not remove that privilege from any user to whom it was granted in a separate `GRANT` statement.

Cascading REVOKEs

When privileges are revoked, the DBMS must decide whether additional revokees are necessary, based on the privileges being revoked. When one revoke causes the DBMS to revoke additional related privileges, it is called *cascading REVOKEs*. Consider the authority hierarchy depicted in [Figure 14.2](#). Joe has been granted the ability to grant a privilege, say *X*, to others. He grants *X* to Pete and Phil with the `GRANT` option. Pete in turn grants *X* to Bruce. Joe also grants *X* to Don, but without the `GRANT` option.

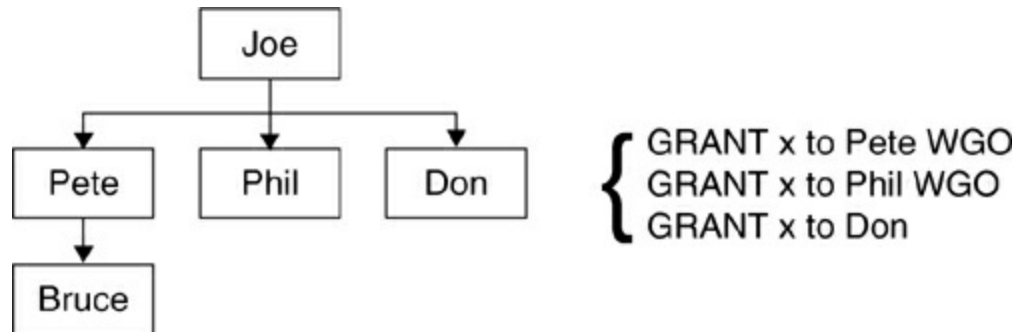


Figure 14.2. Cascading REVOKEs

Now let's investigate the impact of cascading `REVOKEs` by outlining what happens if we revoke the *X* privilege from Joe. In this case, not only will Joe no longer have the *X* privilege, but the DBMS will also remove the authority from Pete, Phil, and Don. Furthermore, because Phil's *X* privilege was revoked, the effect of the revoke will cascade to Bruce, too.

To minimize the impact of cascading `REVOKEs`, avoid granting privileges using the `WITH GRANT OPTION`. The fewer users who can grant subsequent privileges, the easier it is to manage and administer a viable DBMS security infrastructure.

Some DBMS products provide an option of the `REVOKE` statement

enabling authority to be revoked without cascading.

Chronology and Revokes

The timing of a GRANT or REVOKE statement may have a bearing on its impact. For example, in some DBMSs it is possible to grant a privilege to all users except a specific user by issuing the following statements:

```
GRANT DELETE on titles to public;  
COMMIT;  
REVOKE DELETE on titles from userx;
```

The timing of a GRANT or REVOKE statement may have a bearing on its impact.

The first statement grants everyone the authority to delete data from the titles table. Because the REVOKE statement is issued after the GRANT to public, the individual userx is barred from deleting data from this table. Some DBMSs (e.g., DB2) do not permit such exclusions, because the PUBLIC authority will override any revokes; other DBMSs (e.g., Microsoft SQL Server) do permit such exclusions.

The DBA needs to understand exactly how GRANTs and REVOKEs work for each DBMS being administered in order to properly manage privileges on database objects and resources.

Label-Based Access Control

As database systems and applications become more sophisticated, the need for low-level access control to the data becomes more critical. A growing number of DBMSs offer label-based access control (LBAC) to ensure that each piece of data is secured such that only authorized users can perform authorized functions. With LBAC it is possible to support applications that need a more granular security scheme. LBAC can be set up to specify who can read and modify data in individual rows and/or columns.

LBAC provides for a more granular security scheme, specifying who can read and modify data in individual rows and/or columns.

For example, you might want to set up an authorization scenario such that employees can see their own data but no one else's. Or, your authorization

needs might dictate that each employee's immediate manager is able to see his payroll information as well as all of his employees' data, and so on up through the org chart. Setting up such a security scheme is virtually impossible without LBAC.

An administrator must configure the LBAC system by creating security label components, which are database objects used to represent the conditions determining whether a user can access a piece of data. A security policy, composed of one or more security label components, is used to describe the criteria for determining who has access to what data. The security administrator defines the policy by creating security labels that are composed of security label components.

Once created, a security label can be associated with individual columns and rows in a table to protect the data held there. An administrator allows users access to protected data by granting them security labels. When a user tries to access protected data, that user's security label is compared to the security label protecting the data. The protecting label will block some security labels but not others.

Any individual user can hold security labels for multiple security policies. For any given security policy, however, a user can hold at most one label for read access and one label for write access. A security administrator can also grant exemptions to users. An exemption allows users to access protected data that their security labels might otherwise prevent them from accessing. Together, security labels and exemptions are called LBAC credentials.

Any attempted access to a protected column when the LBAC credentials do not permit that access will fail. If users try to read protected rows not allowed by their LBAC credentials, the DBMS simply acts as if those rows do not exist. Even aggregate functions, such as COUNT(*), will ignore rows when the LBAC credentials do not allow read access. This is important because sometimes even having knowledge that the data exists (without being able to access it) must be protected.

To use LBAC, you will need to add a specially named column to act as the security label. The security label column is matched with the multilevel security hierarchy set up by the administrator. For example, you might want to set up a hierarchy using colors, as shown in [Figure 14.3](#). Of course, the hierarchy need not be so complex—you might choose to use something simpler, such as TOP SECRET, SECRET, and UNCLASSIFIED.

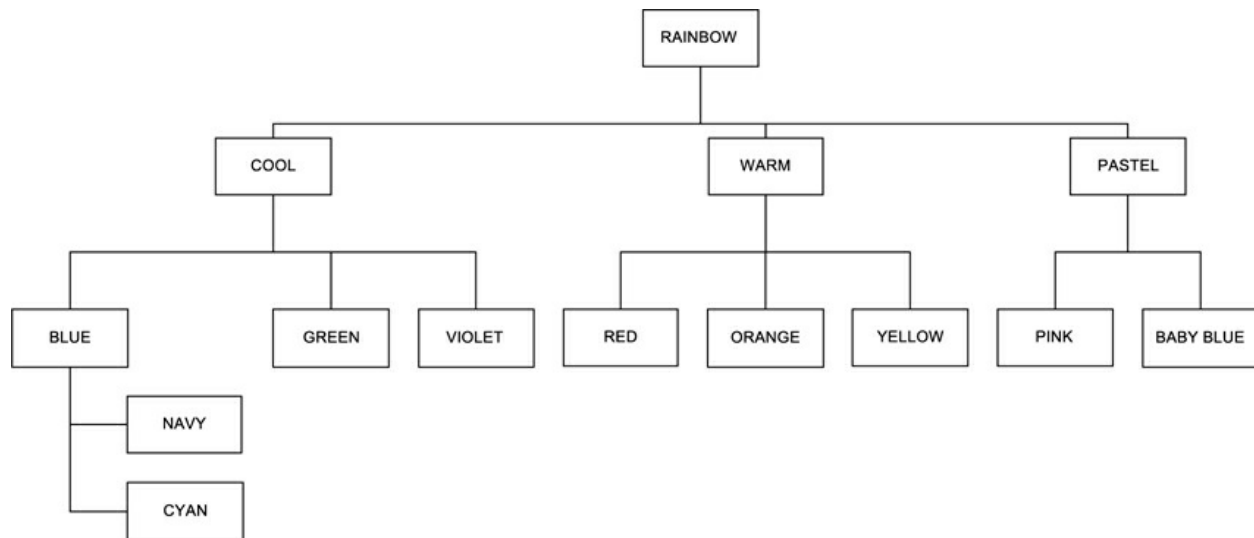


Figure 14.3. Sample LBAC security hierarchy

Consult your DBMS documentation for where and how to establish this hierarchy. For example, with DB2 for z/OS the hierarchy is established in security software, such as RACF. At the top of the hierarchy in the example ([Figure 14.3](#)), the RAINBOW security label includes everything—all colors are included in the rainbow. Middle levels of the hierarchy that represent additional security groupings can be created. In this case we have a second level of COOL, WARM, and PASTEL. Under COOL, we see the cool colors, BLUE, GREEN, and VIOLET. Even further categorization is possible as shown with the various shades of BLUE. And so on throughout the hierarchy. Such a hierarchy provides great flexibility for assigning various levels of security to data in your tables—and thereby appropriately securing the data from unauthorized user access.

Assigning a security label column in your table, and populating that column with the appropriate specification from the hierarchy, makes it possible to control whether or not an individual user can access that particular row. So assigning a row to the security label of RAINBOW makes the data accessible to anyone, whereas assigning the row to NAVY would be much more particular and prohibitive. Users with authority to access RAINBOW can access data because all colors fall under the rainbow. BLUE data can access any row associated with BLUE, or any subordinate blue shade (NAVY and CYAN).

Assigning a security label makes it possible to control whether

or not an individual user can access a particular row.

LBAC offers a very powerful security setup. Such strict security controls are common requirements for government databases, for example.

Security Reporting

Once granted, the DBA will need to monitor and report on the privileges held by users. Database security is maintained in the system catalog. The DBA can use SQL to retrieve the needed information from the appropriate system catalog tables. Alternatively, some DBMSs provide views and system-stored procedures that simplify the retrieval of database security.

As a further consideration, though, be sure to adequately protect the security of the system catalog, especially within the production system. Only the DBA, SA, and security administrator require access to the database security information stored in the system catalog.

Be sure to adequately protect the security of the system catalog.

User security requirements and expectations tend to evolve over time. As new applications are added and business requirements change, database security will need to change. Security reviews should be performed on a regular basis to ensure that database security as implemented continues to match the current user requirements. Reports from the system catalog tables can be used to provide the input for such reviews.

Authorization Roles and Groups

In addition to granting privileges to individual users, the DBMS may provide the capability to assign users

- Specific privileges to a role, which is then granted to others, or
- Specific built-in groups of privileges

Of course, the terminology is not strict among the major DBMSs. Some DBMSs refer to roles as groups, and vice versa. As a DBA, you will need to understand how each DBMS you manage implements roles and groups and how each of these features can be used to simplify database security administration.

Roles

Once defined, a [role](#) can be used to grant one or more preassigned privileges to a user. A role is essentially a collection of privileges. The DBA can create a role and assign certain privileges to that role. Then the role can be assigned to one or more users. The administration of database security is simplified in this way. For example, consider the following sequence of statements:

The DBA can create a role and assign certain privileges to that role.

[Click here to view code image](#)

```
CREATE role MANAGER;
COMMIT;
GRANT select, insert, update, delete on employee to
MANAGER;
GRANT select, insert, update, delete on job_title to
MANAGER;
GRANT execute on payroll to MANAGER;
COMMIT;
GRANT MANAGER to user1;
COMMIT;
```

This script creates a new role named MANAGER, grants privileges on certain tables and procedures to the role, and then assigns user1 the MANAGER role. Additional users can be assigned the MANAGER role, and the DBA will not need to remember to issue each of the individual GRANT statements, because they have already been assigned to the MANAGER role.

Groups

Group-level authority is similar to roles. However, each DBMS provides built-in groups that cannot be changed. Each DBMS implements group-level database security in different ways and with different group names and privileges. There are some similarities across DBMSs, though. The following groups are common among the major DBMSs:

- *System administrator*. Sometimes abbreviated SA or SYSADM, the system administrator group is the most powerful within the DBMS. A user granted SA-level authority typically can execute all database commands and access all databases and objects. The system administrator is usually responsible for installing the DBMS and is

viewed as the owner of system resources and system catalog tables.

- *Database administrator*. Sometimes abbreviated as DBADM or DBA, the database administrator group gives all privileges over a specific database, plus the ability to access, but not modify, data in tables within that database. Users assigned DBA-level authority can drop and alter any objects within the database (tablespaces, tables, and indexes).
- *Database maintenance*. Sometimes abbreviated as DBMAINT, the database maintenance group includes the specific database privileges for maintaining database objects (such as the ability to run utilities and issue commands). Like the DBA group, the DBMAINT-level privilege is granted on a database-by-database basis.
- *Security administrator*. The security administrator role has the privilege set permitting the granting and revoking of database security across the DBMS. Any database security-related activities can be performed by the security administrator, including login and password administration, auditing, security configuration, as well as grants and revokes. Another common name for the security administrator role is SSO.
- *Operations control*. Sometimes referred to as OPER or SYSOPR, the operations control role has the authority to perform operational database tasks such as backup and recovery, or terminating runaway tasks.

Each DBMS provides built-in groups that cannot be changed.

Limit the Number of SA Users

A single organization should limit the number of users who are assigned the SA role or group-level authority.¹ A user with SA capabilities is very powerful. Only corporate DBAs and systems programmers should be granted this level of authority. End users, managers, and application development personnel do not need SA authority to do their jobs.

A user with SA capabilities is very powerful.

Group-Level Security and Cascading REVOKEs

Depending on the group, some users who have been assigned group-level privileges can grant privileges to other users. If the group-level authority is revoked from such a user, any privileges that user granted will also be

revoked. This is similar to the cascading REVOKEs that occur as a result of the WITH GRANT option.

Before revoking a group-level authorization from a user, be sure to ascertain the impact of cascading REVOKEs, and be prepared to reapply the required privileges that will be removed due to the cascading effect.

Other Database Security Mechanisms

Modern relational DBMS products support many capabilities and qualities that can aid in securing data. Some of these capabilities are not primarily security features. For example, views and stored procedures can be used for security purposes, even though that is not their main purpose.

Using Views for Security

Most database security is performed using the native security of the DBMS. However, it is possible to simplify some aspects of database security by creating views to protect your data.

Your organization has deployed an employee table that houses pertinent information about all employees. Columns within the table exist to store the employee's first and last name, middle initial, address, telephone number, salary, and so on. Granting the SELECT privilege on the employee table to a group of users can cause a security problem. While application security is maintained with this scenario, personal security is not because the user could access the personal details, including salary information, of fellow employees.

A view could be created that omits the sensitive information from the employee table. When the view is created without the sensitive columns, users can be granted the SELECT privilege on the view and will be able to access employee information that is deemed appropriate. For example:

A view can be created that omits sensitive information.

[Click here to view code image](#)

```
CREATE view emp_all
AS
SELECT first_name, last_name, middle_initial,
       street_address, state, zip_code
FROM   employee;
```

This simple example shows a view that specifies only certain columns from the base table. Once the view has been created and the user has been granted the SELECT privilege on the view, only the information specified in the view can be retrieved. Eliminating columns from a base table to create a view is referred to as *vertical restriction*.

Of course, the definition of *sensitive* will vary from organization to organization. In our example, the salary and telephone information were removed. It is quite simple to understand why salary is sensitive, but what about telephone number? And if telephone number is sensitive, perhaps the employee's address should be as well. Views allow you to easily specify the column-level security deemed necessary for your organization.

Vertical restriction using views is an alternative to specifying columns when granting table privileges. It also can be easier to implement and administer.

Views can also be used to provide row-level security based on the content of data. This is called *horizontal restriction* and is implemented by coding the appropriate WHERE clauses into the view. For example:

[Click here to view code image](#)

```
CREATE view emp_dept20
AS
SELECT first_name, last_name, middle_initial,
       street_address, state, zip_code
FROM   employee
WHERE  deptno = 20;
```

When users select from the view, only rows that match the predicate will be returned. This view will return only those employees who work in department 20. When users modify rows of the view, if the WITH CHECK OPTION has been specified, the predicates will ensure that values cannot be updated or inserted out of the range. Additionally, rows that do not match the predicate cannot be deleted using the view when WITH CHECK OPTION is specified.

Using Stored Procedures for Security

Stored procedures can be used to provide an additional level of security. The privilege to execute a stored procedure must be explicitly granted or revoked, regardless of the security implemented on the underlying base tables.

The privilege to execute a stored procedure must be explicitly granted or revoked.

Stored procedures can be coded that access only row- and/or column-level subsets of data. The ability to execute these stored procedures can then be granted to users. If no privileges on the underlying base tables are granted, users will be able to access the data only by executing the stored procedure, thereby providing the requisite security.

In addition to providing a level of security, this method can provide better performance if the algorithms in the procedure are coded properly.

Logic-Oriented Security

Sometimes it is necessary to implement security based on an algorithm. For example, what happens if only a subset of users can access a specific table during a specific time of day? This criterion can be coded into the stored procedure. Whenever the stored procedure is executed, it checks the user and the time of day before permitting access.

Some DBMS products offer special functionality for database security. See the sidebar “[Oracle Virtual Private Database](#)” for an example of such functionality.

Encryption

Another security feature for data protection is encryption. Data encryption is a process whereby data is transformed using an algorithm to make it unreadable to anyone without the decryption key. The general idea is to make the effort of decrypting so difficult as to outweigh the advantage to a hacker of accessing the unauthorized data. The general flow of encryption and decryption is depicted in [Figure 14.4](#).

Encryption transforms data, rendering it unreadable to anyone without the decryption key.

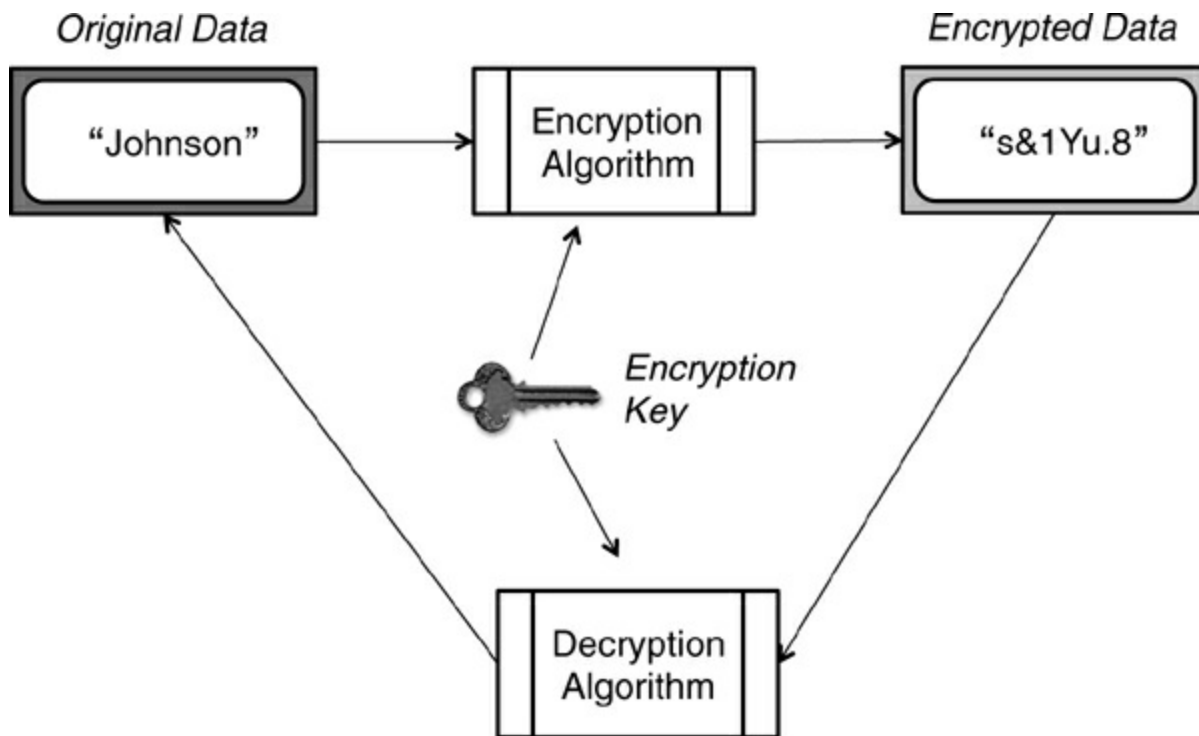


Figure 14.4. Encryption and decryption

Oracle Virtual Private Database

Oracle provides row-level access control through its Virtual Private Database (VPD) technology. VPD is enabled by associating one or more security policies with tables or views. When the table is accessed, either directly or indirectly, the database will consult a function that implements the policy. The policy is basically a SQL predicate (or WHERE clause) that the database appends to the user's SQL statement. This dynamically modifies the user's data access.

With VPD, a user is able to retrieve and manipulate only data that matches the WHERE clause in the policy. In essence, this works something like creating a dynamic view (using WITH CHECK OPTION) that is always applied and enforced.

Encryption has been used by governments and military organizations for years to enable secret data transmission and communication. Because of the added data protection and regulatory requirements of late, encryption is being more commonly used in many business systems and applications.

There are two types of situations where data encryption can be deployed: data in transit and data at rest. In a database context, data “at rest” encryption protects data stored in the database, whereas data “in transit” encryption is used for data being transferred over a network.

Data at Rest Encryption

Encrypting data at rest is undertaken to prohibit “behind the scenes” snooping for information. Consider, for example, a database containing a top-secret military plan. Of course, this data can, and should, be protected using traditional database security and authorization methods, such as granting access to only those with the proper clearance. Additional precaution can be taken by using LBAC techniques if the database contains data needing varying levels of clearance requirements. But what if the data is accessed outside the control of the DBMS?

When the data at rest is encrypted, even if a hacker surreptitiously gains access to the data behind the scenes, without the decryption key the data will be meaningless.

Without the decryption key the data will be meaningless.

Other examples of data at rest encryption in database systems include encrypting backup data and encrypting user IDs and passwords.

Data at rest encryption most commonly is supported by using SQL functions or through an add-on encryption product. Oracle supports transparent data encryption through the use of a “wallet” (see the sidebar).

Data in Transit Encryption

Encrypting data in transit is undertaken to prohibit network packet sniffing. If the data is encrypted before it is sent over the network and decrypted upon receipt at its destination, it is protected along its journey. Anyone nefariously attempting to access the data en route will receive only encrypted data. And again, without the decryption key, the data cannot be deciphered.

Data in transit encryption most commonly is supported using DBMS system parameters and commands or through an add-on encryption product.

Encryption Techniques

There are various encryption techniques and algorithms supported by modern

DBMSs. Among the more common are Data Encryption Standard (DES), Triple DES (TDES), and Advanced Encryption Standard (AES)-128 bit data encryption/decryption, as well as Secure Hash Algorithm (SHA-1) and SHA-256 hashing.

Oracle Transparent Data Encryption

Transparent data encryption enables you to encrypt individual table columns or an entire tablespace. When a user inserts data into an encrypted column, transparent data encryption automatically encrypts the data. When users select the column, the data is automatically decrypted.

Data is encrypted using a wallet, which is an operating system file located outside the database. The database uses the wallet to store the encryption key. The wallet is created using the ALTER SYSTEM command (or within Oracle Enterprise Manager). The wallet itself is encrypted and requires a password as its encryption key. To access the contents of the wallet you must know the password. After the wallet is created, you must open the wallet using the password so that the database can access the master encryption key. The location of the wallet is specified in the sqlnet.ora file.

So, when a user enters data, Oracle retrieves the master key from the wallet, decrypts the encryption key using the master key, uses the encryption key to encrypt the data entered by the user, and stores the encrypted data in the database. When a user accesses data, the process is similar. Oracle decrypts the data and displays it in clear text format.

SQL Injection

Another aspect of database security is designing your applications properly in order to avoid SQL injection attacks. SQL injection is a form of Web hacking whereby SQL statements are specified in the fields of a Web form to cause a poorly designed Web application to dump database content to the attacker. SQL injection is a favorite technique of Web hackers, and stories abound in the news of the technique being used for nefarious purposes (see the sidebar titled "[Examples of SQL Injection Attacks](#)").

SQL injection is a form of Web hacking whereby SQL

statements are specified in a Web form to expose data to the attacker.

In order for SQL injection to succeed, the application software used by the Web site must be vulnerable to an injection attack. SQL injection relies upon programs that do not adequately filter for string literal escape characters embedded in SQL statements or where user input is not strongly typed. So instead of data being input into a form, SQL statements are supplied. The SQL is “injected” from the Web form into the database, causing it to be executed and to access (or even modify) unintended data.

Perhaps it is easiest to comprehend SQL injection by example. Consider a Web-based application using dynamic SQL. The Web site requires users to log in with their e-mail address and a password. Almost all sites of this type also offer an option to retrieve your password by supplying your e-mail address. Perhaps the SQL looks something like this:

```
SELECT userid, password
FROM uid_pwd_table
WHERE field = '$EMAIL';
```

Examples of SQL Injection Attacks

There are many examples of SQL injection being used to attack popular Web sites. For example:

On January 13, 2006, Russian hackers broke into a Rhode Island government Web site and obtained credit card data from residents who had done business online with state agencies (www.xiom.com/whid-2006-3).

On June 29, 2007, a hacker defaced the Microsoft Web site in the United Kingdom using a SQL injection attack (www.cgisecurity.com/2007/06/hacker-defaces.html).

In January 2008, tens of thousands of PCs were infected by an automated SQL injection attack that exploited a vulnerability in application code that accesses Microsoft SQL Server databases (www.pcworld.com/businesscenter/article/146048/mass_sql_injection_atta

On September 19, 2010, during the Swedish election, a voter attempted an injection by hand-writing SQL statements as part of a write-in vote (<http://alicebobandmallory.com/articles/2010/09/23/did->

[little-bobby-tables-migrate-to-sweden](#)).

In June 2011, the Public Broadcasting System (PBS) was hacked, through use of SQL injection (<http://blog.imperva.com/2011/05/pbs-breached-how-hackers-probably-did-it.html>).

In September 2011, Turkish hackers accessed NetNames DNS records and changed entries redirecting users (of the *Telegraph*, the *Register*, the *National Geographic*, and others) to a site set up by them (www.zone-h.org/news/id/4741).

Other Web sites attacked by SQL injection in 2011 included Barracuda Networks, mysql.com, Lady Gaga's Web site, Nokia's developer site, and Canon, among other prominent organizations.

The variable \$EMAIL represents the input from the form on the Web site. A savvy hacker can attempt a SQL injection attack by entering

```
anything' OR '1'='1
```

If the application does not check the input properly, the injection causes the SQL to now look like this:

```
SELECT userid, password
FROM uid_pwd_table
WHERE field = 'anything' OR '1'='1';
```

This will cause a complete dump of every user ID and password in the database because the OR '1'='1' component will always evaluate to TRUE. It does not matter what the first part of the injection was; it could be anything, because the second part of the injection gives the hacker everything in the table.

Another form of SQL injection relies upon improper typing, for example, not checking whether data that should be numeric is actually numeric. Consider, for example:

[Click here to view code image](#)

```
statement := "SELECT * FROM userinfo WHERE id = " +
in_var + ";"
```

In this case, the SQL is being built into the statement variable; the in_var is the variable used to supply the input. Let's assume that the id column is

numeric. However, if the program does not check the data type of the `in_var` variable to ensure that numeric data is supplied, SQL injection can occur. For example, instead of just supplying a numeric value, the hacker can supply something like this:

```
4;DROP TABLE customer
```

If this SQL statement is built and executed, the customer table (if one exists) will be dropped from the database.

SQL Injection Prevention

Using well-designed query language interpreters and coding applications appropriately can prevent SQL injection attacks. When possible, use static SQL (see the next section for more details).

Always validate user input by testing type, length, format, and range. The program should make absolutely no assumptions about the size, type, or content of the data that is received. Test the size and data type of input and enforce appropriate limits. Doing so can help to prevent buffer overruns. Test the content of string variables and allow only expected values to be processed. Any input that contains binary data, escape sequences, and comment characters should be summarily rejected.

Always validate user input by testing type, length, format, and range.

Avoid concatenating user input that has not been validated. String concatenation is the primary point of entry for SQL injection attacks. Furthermore, consider using stored procedures to validate user input.

Analyze input and reject anything that contains special characters such as the semicolon (;), the string delimiter ('), comment delimiters (--, /*...*/), V\$ (the beginning of Oracle DBA views), and xp_ (the beginning of SQL Server catalog stored procedures).

For Microsoft SQL Server users, beware of the Transact-SQL procedure `xp_cmdshell`. This procedure spawns a Windows command shell and passes in a string for execution. So it is entirely possible for a SQL injection attack not only to steal or damage data, but to upload and execute its own code, probe the network, and even launch attacks against other sites from the victim of the SQL injection. Use of `xp_cmdshell` is disabled by default, and it is a

good idea to keep it that way to thwart malicious users from attempting to elevate their privileges using the command.

With foreknowledge of SQL injection techniques and proper development procedures, all SQL injection attacks can be prevented.

Static versus Dynamic SQL

Using static SQL instead of dynamic SQL can improve the security of your database applications and data. Static SQL is hard-coded into the application and cannot be changed at run time. Dynamic SQL is flexible and can change at run time. When the SQL can change at run time, a sufficiently motivated and skilled hacker can potentially change the SQL or potentially deploy a SQL injection attack to gain access to unauthorized data.

Using static SQL instead of dynamic SQL can improve security.

Static SQL is common in mainframe DB2 applications, but not so much for other platforms and database systems. The DB2 BIND command “hardens” the SQL and optimizes access to the data. See the sidebar on the DB2 BIND command for additional details.

If you are using a CLI, such as ODBC or JDBC, to access the database, static SQL is probably not an option. However, you can use stored procedures to gain a measure of protection similar to that of static SQL.

The DB2 BIND Command

The BIND command can be thought of as a compiler for SQL statements. In general, the BIND command reads SQL statements embedded in the application program and produces an executable access path to the data as directed by the SQL statements being bound. Once the SQL is statically bound, it cannot be changed without undergoing the BIND process again.

Binding performs many functions to create packages that access the requested data, including reading the SQL statements in the program, syntax checking, metadata validation, authorization validation, and optimization.

An added security feature of DB2 binding is provided with the ENABLE and DISABLE options, which can be specified to control the environment in which the plan or package being bound can be executed.

ENABLE ensures that the program operates in only the enabled environments. DISABLE permits execution of programs everywhere except for those environments explicitly disabled. ENABLE and DISABLE are mutually exclusive parameters; only one can be used per program package. For example, if a program is bound specifying ENABLE(CICS), only the CICS transaction manager is permitted to execute the program. Binding in this manner enables DBAs to effectively prohibit programs that are designed to be executed in one environment (e.g., online CICS) from being surreptitiously used in another (e.g., batch).

Auditing

Auditing is a DBMS facility that enables DBAs to track the use of database resources and privileges. When auditing is enabled, the DBMS will produce an audit trail of database operations. Each audited database operation produces an audit trail of information, including what database object was impacted, who performed the operation, and when. Depending on the level of auditing supported by the DBMS, an actual record of what data actually changed may also be recorded. Tracking who does what to which data when is important because there are many threats to the security of your data. (See the “[Threats to Security](#)” sidebar.)

Auditing enables DBAs to track the use of database resources and privileges.

Keep in mind that auditing tracks what a particular user has done once access has been allowed. Auditing occurs post activity; it does not do anything to prohibit access. Audit trails help promote data integrity by enabling the detection of *security breaches*, also referred to as *intrusion detection*. An audited system can serve as a deterrent against users tampering with data because it helps to identify infiltrators.

Auditing occurs post activity; it does not do anything to prohibit access.

Threats to Security

External agents trying to compromise your security and access your company data are rightly viewed as a threat to security. However, industry studies have shown that many security threats are internal—within your organization. The most typical security threat is a disgruntled or malevolent current or ex-employee who has valid access to the DBMS. Auditing is crucial because you may need to find an instance of unauthorized access by an authorized user.

Database auditing is covered in more detail in [Chapter 15, “Regulatory Compliance and Database Administration.”](#)

External Security

In addition to database security, the DBA must ensure that certain resources used by the DBMS are protected from being accessed outside the control of the DBMS. If database resources are not accessed using DBMS commands and SQL statements, database security mechanisms cannot be relied on to enforce proper user authentication.

When using external security mechanisms to protect database-related resources, the DBA should focus primarily on the data sets and files used by the DBMS. Data sets to protect at the operating system or file system level include

- System catalog data files
- Active and archive log files
- User data sets for tablespaces
- User data sets for indexes
- Audit data files
- Performance trace files
- Program and script files (both source and executable code)

Focus primarily on the data sets and files used by the DBMS.

Ingenuous users intent on mischief may be able to figure out the format of these files and access unauthorized data if you fail to protect these data sets and files. An additional level of protection can be achieved by compressing the data within the DBMS. This places the additional burden on the hacker of

trying to decompress the data. Of course, compression is not sufficient protection.

If data encryption software is available for use within database files, consider using it. Data encryption is a security technique that encodes legible data into a scrambled format, making the files unreadable without the encryption key. The general idea is to make the effort of decrypting so difficult as to outweigh the advantage to the hacker of accessing the unauthorized data.

Additional security may need to be applied to DBMS system resources, such as the physical storage and address spaces used to run the DBMS, the DBMS console, and files used to install the DBMS.

Job Scheduling and Security

Most organizations schedule tasks to be run at predetermined times, and when those tasks involve database access, authority must be granted to the scheduler. Scheduling is usually accomplished using a third-party job scheduler such as CA-7, Control-M, or AutoSys. When scheduling software is used to control the submission and scheduling of batch programs and scripts, the DBA will have to determine the best way to grant database security to the scheduler.

It is not a very good idea to grant SYSADM authority to the job scheduler. Doing so would permit any job to perform any database task—creating potentially severe security problems. Instead, determine how to grant individual authorization to specific jobs using the facilities of the scheduling package and the DBMS. Many job schedulers can be set up to generate a user ID for each job. The generated ID can be granted the proper authority based on the types of actions that are authorized for that particular job.

It is not a good idea to grant SYSADM authority to the job scheduler.

Another common security mistake made at some shops is embedding actual passwords into database utility jobs and scripts. If the password is hard-coded into the job, anyone can read it and use it elsewhere in the system. This does not protect the security of your data.

Non-DBMS DBA Security

The DBA will need to possess a fairly high level of operating system authority in order to perform the job of administering and managing the organization's databases and data. For example, in the UNIX environment some installation tasks require root authority. This situation can be handled in two ways: Either grant the DBA root authority to do the installation or turn the specific installation tasks requiring root authority over to the UNIX system administrator. Either option is viable. My preference is to grant the authority to the DBAs if the DBA staff possesses the requisite level of UNIX skills to understand the ramifications of having root authority. Either way, though, the DBAs and SAs will need to cooperate in order to create an effective operating system security approach that enables the DBAs to perform their job while at the same time protecting the security and integrity of the platform.

DBMS Fixpacks and Maintenance

Another important consideration with regard to database security is the development and implementation of a DBMS maintenance plan. Like any system software, the DBMS itself must be periodically patched to remedy security defects and other software bugs in the DBMS code. Failing to apply patches as they are delivered from the DBMS vendor via fixpacks, downloads, and other maintenance streams can cause security issues in your databases.

Refer to the sidebar on Oracle Database patches for a discussion of how one vendor, Oracle, manages the software maintenance of its DBMS.

Oracle Database Patches

Oracle has implemented a consistent process for patching its system software, including its DBMS. There are two types of patches for which DBAs need to be prepared: Critical Patch Updates and Security Alerts.

Critical Patch Updates are the primary means of releasing security fixes for Oracle products to customers with valid support contracts. They are released on the Tuesday closest to the seventeenth day of January, April, July, and October.

Security Alerts are issued by Oracle for vulnerability fixes deemed to be too critical to wait for the next Critical Patch Update.

Of course, you should evaluate each patch on a case-by-case basis. You may want to apply critical security patches as soon as they are available. On the other hand, you may want to delay applying patches that are not deemed to be urgent by the DBMS vendor. Remember, when you apply a patch, you are changing your system software and you could be introducing additional flaws and defects. Unless you have the time and resources to conduct a thorough and rigorous test of every patch, only urgent or critical patches should be applied quickly.

Some fixpacks contain literally hundreds of software patches. Although it is a good idea to strive for currency in your DBMS software, use prudent judgment before applying every patch the DBMS vendor supplies to you.

Summary

Database security is an important component of a DBA's job. Without a comprehensive database security plan and implementation, the integrity of your organization's databases will become compromised. DBAs should learn the security mechanisms at their disposal to assure that only authorized users are accessing and changing data in the company's databases.

DBAs should learn the security mechanisms at their disposal.

Furthermore, DBAs should implement auditing operations to verify that the database security measures being deployed are adequate.

Review

1. What is the purpose of a login?
2. Explain two ways that views can be used to implement data security.
3. Once a password has been established for a login, it cannot be changed: true or false?
4. What two SQL DCL statements are used to establish and remove database privileges?
5. If a password needs to be deactivated for a period of time, but may need to be reactivated later, it can be _____ instead of _____.
6. What is a cascading REVOKE and what effect can it have on database security?

7. Is “2926Glenmore” a good password? Why or why not?
8. What is the effect of granting a privilege to PUBLIC?
9. Why is it important to protect some database resources using security mechanisms external to the DBMS?
10. Explain four useful ways of helping to prevent SQL injection attacks.

Suggested Reading

Afyouni, Hassan A. *Database Security and Auditing: Protecting Data Integrity and Accessibility*. Boston, MA: Thomson (2006). ISBN 0-619-21559-3.

Ben Natan, Ron. *Implementing Database Security and Auditing*. Oxford, UK: Elsevier Digital Press (2005). ISBN 978-1-55558-334-7.

Bond, Rebecca, et al. *Understanding DB2 9 Security*. Indianapolis, IN: IBM Press (2007). ISBN 0-13-134590-7.

Castano, Silvana, Mariagrazia Fugini, Giancarlo Martella, and Pierangela Samarati. *Database Security*. Wokingham, England: Addison-Wesley/ACM Press (1995). ISBN 0-201-59375-0.

Clarke, Justin. *SQL Injection Attacks and Defense*. Burlington, MA: Syngress (2009). ISBN 978-1-59749-424-3.

Denning, Dorothy E. *Cryptography and Data Security*. Reading, MA: Addison-Wesley (1982). ISBN 0-201-10150-5.

Ingram, Aaron, and Josh Shaul. *Practical Oracle Security*. Rockland, MA: Syngress (2007). ISBN 978-1-59749-198-3.

Kenan, Kevin. *Cryptography in the Database: The Last Line of Defense*. Boston, MA: Symantec Press/Addison-Wesley (2006). ISBN 0-321-32073-5.

Litchfield, David, et al. *The Database Hacker’s Handbook: Defending Database Servers*. Indianapolis, IN: John Wiley & Sons (2005). ISBN 0-7645-7801-4.

Oram, Andy, and John Viega. *Beautiful Security*. Sebastopol, CA: O’Reilly (2009). ISBN 978-0-596-52748-8.

Thuraisingham, Bhavani. *Database and Applications Security: Integrating Information Security and Data Management*. Boca Raton, FL: Auerbach

(2005). ISBN 0-8493-2224-3.

15. Regulatory Compliance and Database Administration

If you've been head-down in the trenches, you might have missed the sweeping changes being thrust upon the data world due to regulatory compliance. And even if you've noticed, chances are that the sheer volume and nature of the regulations were too mind-boggling to fully digest. This chapter provides an introduction to the regulatory landscape and its impact on database administration and management.

There are many industry and governmental regulations driving the need to improve data protection, management, and administration. Let's review a couple of the high-visibility regulations.

One of the more visible governmental regulations is the Sarbanes-Oxley (SOX) Act, officially known as the U.S. Public Company Accounting Reform and Investor Protection Act of 2002. The goal of SOX is to regulate corporations in order to reduce fraud and conflicts of interest, to improve disclosure and financial reporting, and to strengthen confidence in public accounting. But the cost of the law and its impact on IT and database management are significant. Section 404 of SOX specifies that CFOs must do more than simply vouch that their company's finances are accurate; they must guarantee the processes used to add up the numbers. Those processes are typically computer programs that access data in a database, and DBAs create and manage that data as well as many of those processes. The SOX act also requires that financial data be hardened against unauthorized access, invalid transactions, and any other type of modification that might invalidate financial reporting integrity. If the financial data is in the database, the management and administration practices in use for that database system must be shown to comply. So the Sarbanes-Oxley Act requires that additional rigor and structure be infused into database administration practices and procedures. Of course, DBAs managed corporate database systems prior to SOX. The accounting and regulatory demands of SOX require those practices to be strengthened and formalized. Specifically, database security and access control practices may need to be modified. And a more comprehensive approach to the auditing of data access may be required, especially as it

relates to unusual or suspect activity.

Furthermore, the Sarbanes-Oxley Act demands that a company's financial data must be completely recoverable in the event of a logical or physical failure without loss of data that would invalidate the integrity of the financial reports. And that recovery must be possible in a timely manner. Many organizations may have gone for a long time without testing the viability of their backups for database recovery.

Many organizations may have gone for a long time without testing the viability of their backups.

Consider also the Health Insurance Portability and Accountability Act, more commonly referred to simply as HIPAA. This legislation contains language specifying that health care providers must protect individuals' health care information, even going so far as to state that the provider must be able to document everyone who even so much as looked at their information. HIPAA audits frequently require the examination of the processes used to create, document, and review exception reports and logs. When confronted with a HIPAA audit, organizations can be required to produce a list of exceptions to policy, such as "When were patient records accessed during off hours and by whom?" Without database auditing software, it is impossible to produce a list of users who looked at a specific row or set of rows in any database.

The Gramm-Leach-Bliley Act (GLB Act), also known as the Financial Modernization Act of 1999, is a federal law enacted in the United States to control the ways that financial institutions deal with the private information of individuals. The GLB Act consists of three sections:

- The Financial Privacy Rule, which regulates the collection and disclosure of private financial information
- The Safeguards Rule, which stipulates that financial institutions must implement security programs to protect such information
- The Pretexting provisions, which prohibit the practice of pretexting (accessing private information using false pretenses)

The GLB Act also requires financial institutions to give customers written privacy notices that explain their information-sharing practices.

The E-Government Act was passed in 2002 as a response to terrorist

threats. Title III of the act is named the Federal Information Security Management Act (FISMA). FISMA basically states that federal agencies, contractors, and any entity that supports them must maintain security commensurate with potential risk. Officials are graded on the potential effect a security breach would have on their operations.

SOX, HIPAA, GLB, and FISMA are examples of governmental regulations. But there are industry regulations that can be just as daunting in terms of compliance. The most visible industry regulation is certainly PCI DSS, which stands for Payment Card Industry (PCI) Data Security Standard (DSS). It was developed by the major credit card companies to help prevent credit card fraud, hacking, and other security issues. A company that processes, stores, or transmits credit card numbers must be PCI DSS compliant or it risks losing the ability to process credit card payments.

Regulatory compliance holds a powerful sway over upper-level management at most medium-size to large organizations because of its potent impact. Business executives are keenly aware of the need to comply, although they are not always aware of all the details that involves. This is so because failure to comply can result in prosecution, which may involve huge fines and even imprisonment (see the sidebar titled “[The Reality of Prosecution](#)”). Regulatory compliance can impose upon C-level executives the need to be able to prove that corporate data (and therefore database systems) is protected and that processes and procedures enacted upon the data are accurate and required.

Regulatory compliance holds an important sway over upper-level management at most medium-size to large organizations.

The cost of compliance can be significant, but so can the cost of noncompliance, especially when you factor in the potential for fines and lost business. And it is not possible to just ignore the regulations and hope that no one notices.

The Reality of Prosecution

Prosecution for compliance failure is not just an idle threat. Consider these examples of successful prosecutions:

On May 25, 2006, Ken Lay (former CEO of Enron Corporation) was found guilty of ten counts against him. Because each count carried a

maximum five- to ten-year sentence, Lay could have faced 20 to 30 years in prison. However, he died about three and a half months before his scheduled sentencing. Jeff Skilling (another former Enron CEO) was found guilty of 19 out of 28 counts against him, including one count of conspiracy and one count of insider trading. Skilling was sentenced to 24 years and four months in federal prison.

In June 2005, 800,000 WorldCom investors were awarded \$6 billion in settlements; the payouts were funded by the defendants in the case, including investment banks, audit firms, and the former directors of WorldCom. The judge in the case noted that the settlements were “of historic proportions.” Additionally, Bernard Ebbers, former CEO of WorldCom, was convicted of fraud and sentenced to 25 years in prison.

In September 2005, Dennis Kozlowski (former Tyco CEO) and Mark Swartz (former Tyco CFO) were sentenced to eight to 25 years in prison for stealing hundreds of millions of dollars from Tyco. Additionally, Kozlowski had to pay \$70 million in restitution; Swartz, \$35 million.

In June 2005, John Rigas (founder and former Adelphia CEO) was sentenced to 15 years in prison. His son, Tony, was also convicted of bank fraud, securities fraud, and conspiracy; he received a 20-year sentence.

Not every prosecution results in a conviction, though. In March 2005, Richard Scrushy, founder and former CEO of HealthSouth, was acquitted of charges relating to a \$2.7 billion earnings overstatement. Scrushy blamed his subordinates for the fraud.

A Collaborative Approach to Compliance

Ensuring compliance requires a collaborative effort among business users, IT, and the legal department. This can prove to be a challenge because these three disparate groups are quite distinct and rarely communicate collectively. IT talks to legal only when they have to—and that is usually just to get approval on contract language for a software purchase. IT and business communicate regularly (at least they should), but perhaps not as effectively as they might. But all three are required:

Compliance requires a collaborative effort among business

users, IT, and the legal department.

- *Business* must understand the legal requirements imposed on their data and systems as dictated in regulations.
- *Legal* must be involved to interpret the legal language of the regulations and ensure that the business is taking proper steps to protect itself.
- *IT* must be involved to implement the policies and procedures to enact the technology to support the regulatory mandates.

Organizations need to map and categorize their business data in accordance with how each data element is impacted by regulations. They need to be able to answer questions such as “Which data elements are under the control of which regulation? And what does the regulation require in the way we manage that data?”

Once mapped, controls and policies need to be enacted that enforce compliance with the pertinent regulations. This can require better protection and security, enforce longer data retention periods, impose stricter privacy sanctions, mandate improved data quality practices, and so on.

Why Should DBAs Care about Compliance?

Compliance starts with the CEO, but it works its way down into the trenches and impacts database administration. The CEO relies on the CIO to ensure that IT processes are compliant; the CIO relies on the IT managers, one of whom (the DBA manager) controls the database systems; and the DBA manager relies on DBAs to ensure that data is protected and controlled.

The impact of regulatory compliance upon database administration is various. DBAs are not responsible for developing and enforcing compliance, but their job is impacted based upon compliance-related projects and responsibilities. The primary impact of compliance on the DBA is in investigating, installing, and managing the technology that supports compliance, particularly regarding data and the DBMS.

The primary impact of compliance on the DBA is in investigating, installing, and managing the technology that supports compliance.

Compliance-related tasks that impact database administration include

- Metadata management and data quality
- Database and data access auditing
- Data masking and obfuscation
- Long-term data retention and database archiving
- Closer tracking of traditional DBA tasks (e.g., change management, backup and recovery)

The remainder of the chapter will delve into these subjects in sufficient depth to educate DBAs on the impact of compliance on their job.

Metadata Management, Data Quality, and Data Governance

Ensuring data quality is a big part of regulatory compliance. Its importance cannot be understated. Poor data quality costs the typical company at least 10 percent of revenue; 20 percent is probably a better estimate.¹ According to software marketing and technology expert Hollis Tibbetts,² “Incorrect, inconsistent, fraudulent and redundant data cost the U.S. economy over \$3 trillion a year.”

Ensuring data quality is a big part of regulatory compliance.

The cost of poor data quality notwithstanding, high-quality data is crucial for complying with regulations. Think about it. If the data is not accurate, how can you be sure that the proper controls are being applied to the right pieces of data to comply with the appropriate regulations?

Metadata

Good data quality starts with metadata. [Chapter 22](#) of this book is dedicated to metadata management, so it will not be covered here. Suffice it to say, accurate data definitions are required in order to apply the controls for compliance to the correct data.

What Is Metadata?

Metadata characterizes data, providing documentation such that data can be understood and more readily consumed by an organization. Metadata answers the who, what, when, where, why, and how questions for users of the data.

Metadata is required to place the data into proper categories for determining which regulations apply. For example, SOX applies to financial data, HIPAA applies to health care data, and so on. Some data will apply to multiple regulations and some data will not be regulated at all. Without proper metadata definitions, it is impossible to apply regulatory compliance to data.

Data Quality

The next step is to ensure that the data, once accurately defined, is itself accurate. Imposing regulatory controls on the wrong data does no good at all. This raises the question “How good is your data quality?” Estimates show that, on average, data quality is an overarching industry problem. According to data quality expert Thomas C. Redman,³ payroll record changes have a 1 percent error rate; billing records have a 2 to 7 percent error rate; and the error rate for credit records is as high as 30 percent.

But what can a DBA do about poor-quality data? Data quality is a business responsibility, but the DBA can help by instating technology controls. Building constraints into the database can improve overall data quality, as well as defining referential integrity in the database as defined in [Chapter 4](#). Additional constraints should be defined in the database as appropriate to control uniqueness, as well as data value ranges using CHECK constraints and triggers.

Data profiling can discover the quality, characteristics, and potential problems of information.

Another technology tactic that can be deployed to improve data quality is data profiling. *Data profiling* is the process of examining the existing data in the database and collecting statistics and other information about that data. With data profiling, you can discover the quality, characteristics, and potential problems of information. Using the statistics collected by the data profiling solution, business analysts can undertake projects to clean up problematic data in the database.

Data profiling can dramatically reduce the time and resources required to find problematic data. Furthermore, it allows business analysts and data stewards to have more control of the maintenance and management of

enterprise data.

Data Governance

Data governance programs are becoming more popular as corporations work to comply with more and stricter governmental regulations. A data governance program oversees the management of the availability, usability, integrity, and security of enterprise data. A sound data governance program includes a governing body or council, a defined set of procedures, and a plan to execute those procedures.

So an organization with a strong data governance practice will have better control over its information. When data management is instituted as an officially sanctioned mandate of an organization, data is treated as an asset. That means data elements are defined in business terms; data stewards are assigned; data is modeled and analyzed; metadata is defined, captured, and managed; and data is archived for long-term data retention.

All of this should be good news to data professionals who have wanted to better define and use data within their organizations. The laws are finally catching up with what we knew our companies should have been doing all along.

Database Auditing and Data Access Tracking

In a world replete with regulations and threats, organizations today have to go well beyond just securing their data. Protecting this most valuable asset means that companies have to perpetually monitor their systems in order to know who did exactly what, when, and how to their data.

Increased regulations demand the implementation of policies and procedures to protect sensitive data. For example, medical professionals must comply with HIPAA, thereby making them ever vigilant in the techniques used to manage and protect the data under their care. When sensitive data is stored in a DBMS, extra attention must be paid to how data in those databases is governed.

Consider also the impact on organizations that accept credit and payment cards. Such companies must comply with PCI DSS to maintain the right to process credit card transactions. Thus, executives must ensure their databases are protected such that only properly authorized entities have access to only the specific data they need in order to do their jobs—and to be able to prove

this.

Database auditing is the process of monitoring access to and modification of selected database objects and resources.

One important technique is to audit access to database data. *Database auditing* is the process of monitoring access to and modification of selected database objects and resources within operational databases and retaining a detailed record of the access that can be used to proactively trigger actions and can be retrieved and analyzed as needed. This can be accomplished using a database auditing facility. [Table 15.1](#) offers an overview of the auditing requirements of several representative regulations.

Table 15.1. Database Auditing Requirements of Several Regulations

Audit Requirement	SOX	PCI DSS	GLB	HIPAA
Access to sensitive data (SELECT)		X	X	X
Modification of sensitive data (INSERT, UPDATE, DELETE)	X			
Database changes/DDDL (CREATE, ALTER, DROP)	X	X	X	X
Security authorizations/DCL (GRANT, REVOKE)	X	X	X	X
Security exceptions (e.g., failed logins)	X	X	X	X

Sensitive corporate data cannot be fully protected by simply setting up database authorization using the controls within the database software. This is so because it is not possible to guarantee that surreptitious access to sensitive data is blocked with simple database authorization mechanisms. And secondarily, it is possible for authorized users to nefariously access data. Database auditing can help protect data in both of these situations.

All of the major DBMS products offer built-in capabilities for auditing databases, but ISVs offer more capable software with more flexible capture technology, prepackaged compliance reports, and multi-DBMS support. Robust database auditing software can comprehensively track the usage of database resources and authority. When auditing is enabled, each database operation produces a detailed audit trail of information, tracking what data was accessed, who accessed it, and when. Operators can analyze the audit trail and generate reports showing access and modification patterns against the sensitive data in the DBMS.

An audit trail can be useful in many situations. Your company's business

practices and security policies may dictate a comprehensive ability to trace every data change back to the initiating user. Any number of regulations may require your organization to analyze data access and produce regular reports. You may be required to produce detailed reports on an ongoing basis, or perhaps you just need the ability to identify the root cause of data integrity problems on a case-by-case basis. Auditing is beneficial for all of these purposes.

Database auditing helps answer questions such as “Who accessed the payment account details for Mr. Jones?” or “When was Mrs. Smith’s appointment time changed?” as well as “Who changed that appointment time?” It is even possible to answer more detailed questions such as “What was the old appointment time prior to the change?” The ability to answer such questions is very important for regulatory compliance.

A typical auditing facility permits auditing at different levels within the DBMS—for example, at the database, database object, and user levels. One of the biggest problems with DBMS audit facilities is performance degradation. The audit trails that are produced must be detailed enough to capture “before” and “after” images of database changes. However, capturing so much information, particularly in a busy system, can cause performance to suffer. Furthermore, this audit trail must be stored somewhere, which is problematic when a massive number of changes occur. Therefore, most auditing facilities allow for the selective creation of audit records to minimize performance and storage problems.

Most auditing facilities allow for the selective creation of audit records.

Although each DBMS offers different auditing capabilities, some common items that can be audited by DBMS audit facilities include

- Login and logoff attempts (both successful and unsuccessful)
- Database server restarts
- Commands issued by users with system administrator privileges
- Attempted integrity violations (where changed or inserted data does not match a referential, unique, or check constraint)
- SELECT, INSERT, UPDATE, and DELETE operations
- Stored procedure executions

- Unsuccessful attempts to access a database or a table (authorization failures)
- Changes to system catalog tables
- Row-level operations

As an additional note, some auditing features and tools can also be used for data recovery. We will examine this aspect of auditing in [Chapter 16](#), “[Database Backup and Recovery](#).”

Auditing can also be used for data recovery.

If you have turned on database auditing at your site, consider the following advice:

- *Auditing can be a large consumer of system resources.* When the audit queue is full, tasks that generate audit records will wait until the auditing task can resume. Consider using a larger audit queue if performance suffers. As a last resort, discontinue auditing when performance is unacceptable.
- *Place the system catalog tables that store security-related information on a separate, inactive disk.* This will enhance auditing performance by decreasing head contention.
- *Ensure that the data set or table used to store audit data does not fill up.* When the audit data set is full, auditing will be disabled, records in the current audit queue will be lost, and any user task attempting to send data to the audit queue will be canceled.

Database Auditing Techniques

There are several popular techniques that can be deployed to audit database data. By far, the best technique engages proactive monitoring of database operations directly at the database server. This technique captures *all* requests for data as they are made. When the audit details are captured at the server level, the software can guarantee that all access is monitored.

The method by which the audit trail is produced is a significant consideration. There are basically five options: deploying native DBMS traces, scanning database transaction logs, adding manual “audit” columns, sniffing network traffic, or tapping requests at the source (the database server). Of these methods, only the last provides 100 percent visibility into

database activities without becoming a performance drain.

There are five techniques for implementing database auditing: traces, log scanning, audit columns, network packet sniffing, and tapping the database server.

Techniques such as trace-based auditing or parsing database logs can miss certain types of database activities. For example, native DB2 auditing is implemented using trace-based auditing. But if a transaction accesses a table more than once in a single unit of recovery, the audit trace records only the first access. This can be unacceptable for some types of regulatory compliance.

Log-based auditing is problematic, too. Because such auditing relies on the database log, the only activities that can be captured are updates, inserts, and deletes. Some regulations (HIPAA, for example) require tracking reads, which never show up on a transaction log.

Some penny-pinching organizations attempt to implement auditing simply by adding “audit” columns to tables, such as `last_modified_date`. The idea behind this approach is that each modifier must also update the `last_modified_date` column whenever any other type of modification is made. But this approach is ill advised, and auditors tend not to like it because

- Audit trails should be kept outside of the database
- If you delete the row, you lose the audit data
- You cannot guarantee that the last modified data is completely accurate because the approach relies on code that may be buggy, or an update program might simply be overlooked and not modify the `last_modified_date` column
- Someone may set the audit column by accident (or nefariously)

This leaves two final approaches: network sniffing and tapping requests at the database server. Both are reasonable approaches, but even network sniffing can miss activity that should be audited. Consider DBA activity performed directly on the database server. If the activity does not go across the network, a network packet sniffer cannot track the actions. Another consideration, especially for mainframe database auditing, is that many mainframe applications run completely within the mainframe without requiring networked requests (for example, a CICS/DB2 application).

Tapping requests directly on the server is the only foolproof method of database activity monitoring that guarantees that all auditable actions are trapped regardless of platform, DBMS, and activity.

Before choosing and implementing a database auditing solution for your database systems, you should investigate the methodology being used. The database auditing solution should offer a surveillance mechanism that is selective, comprehensive, and noninvasive.

A database auditing solution should offer selective, comprehensive, and noninvasive server surveillance.

Selective means that the solution should be rules based to enable the capture of audit details only on the specific data that requires auditing. The general rule of thumb is that only data that must be audited to be in compliance should be audited, and nothing more.

Comprehensive means that the solution must be able to capture the complete scenario of auditable information. It should be able to capture both database reads and modifications. Furthermore, it should not be possible to bypass the audit collection.

And *noninvasive* means that the database auditing software must be able to audit access to data without incurring expensive performance degradation. Database auditing adds overhead, but the last thing you want is for the lights to dim when you turn on database auditing. Be sure to test the efficiency of the audit collection mechanism on production workloads before adopting it.

Furthermore, a robust database access auditing solution that addresses regulatory compliance should be able to provide answers to at least the following questions:

- Who accessed the data?
- At what date and time was the access?
- What program or client software was used to access the data?
- From what location was the request issued?
- What SQL was issued to access the data?
- Was the request successful; and if so, how many rows of data were retrieved?
- If the request was a modification, what data was changed? (A “before”

and “after” image of the change should be accessible.)

Ensuring compliance with tedious government and industry regulations is a daunting task. This, along with the growing need to protect databases from the increasing online and internal threats to sensitive data, has resulted in financial executives being asked to be more personally responsible for the safety of corporate data. Data access auditing solutions can help organizations remain vigilant as they adapt to meet the growing mountain of regulatory requirements safely and proactively.

Privileged User Auditing

One of the most common forms of database auditing is privileged user auditing. DBAs and system administrators typically are granted high-level authority such as DBADM or SYSADM privileges in the database. Users with such authority have carte blanche access to the database instance and all its data. DBAs are trusted agents and should not abuse the overarching privileges they are granted. But the general maxim of “trust, but verify” applies in this case. DBAs need a high degree of authorization to do their job, but that also brings the opportunity for nefarious activity. Implementing privileged user auditing to track every action taken by such users is a wise course of action (see the sidebar “[DBA Steals Data](#)”).

Trust, but verify.

DBA Steals Data

In July 2007, *Computerworld* reported⁴ that the senior DBA at Certegy Check Services, Inc., a subsidiary of Fidelity National Information Services, Inc., stole data belonging to about 2.3 million consumers and sold it to a data broker. The broker turned around and sold some of the data to other companies.

The theft was uncovered when one of Certegy’s customers alerted the company to a correlation between a small number of check transactions and the receipt by the customer of direct telephone solicitations and mailed marketing materials. Certegy launched an immediate investigation and was unable to detect any breach of its security systems and, thereafter, engaged a forensic investigator to validate its findings.

Eventually the culprit was tracked down by following the

relationships between the marketing companies that sold the data. The employee was a senior-level DBA who was entrusted with defining and enforcing data access rights.

By using a database auditing solution to enforce privileged user tracking, you can ensure that trusted users are acting appropriately. And that will make both your privileged users and your internal auditors happy. DBAs get to keep the authority they need to do their job, and the auditor gets to monitor the DBA's activities to ensure that those actions are appropriate.

Data Masking and Obfuscation

Data masking is the process of protecting sensitive information in non-production databases from inappropriate visibility. It ensures that sensitive data is replaced with realistic but not real data. The goal is that sensitive personally identifiable information (PII) is not available outside of the authorized environment. Protecting sensitive data from prying eyes can prevent fraud, identity theft, and other types of criminal activities.

PCI DSS is perhaps the largest regulatory driver of data masking in database systems because it requires that the payment card number be obscured on printed receipts. The next time you use a credit card to make a purchase look at the receipt. The credit card number should not be visible.

Data masking can be done while provisioning test environments so that copies created to support application development and testing do not expose sensitive information and thus avoid risks of leaking. Valid production data is replaced with usable, referentially intact, but incorrect or obfuscated data. After masking, the test data is usable just like production data, but the information content is secure. Another form of data masking applies directly to production data (or production reports), such as in the earlier PCI DSS example. In other words, the data is accurate in the database but masked upon retrieval and display.

What Is PII?

Personally Identifiable Information (PII) is information that can be used to uniquely identify, contact, or locate a person or can be used with other sources to uniquely identify an individual. Many regulations stipulate the manner in which PII must be handled and protected.

Examples of PII include full name, national identification number (such as a Social Security number in the United States), IP address, vehicle identification number (VIN), fingerprints, date of birth, birthplace, credit card numbers, driver's license number, and genetic information.

Data Masking Techniques

It is possible to mask data using a variety of techniques. A good data masking solution should offer the ability to mask using multiple techniques. Common techniques include the following:

Data masking solutions should be able to mask using multiple techniques.

- The *substitution* technique replaces existing data with random values from a preprepared data set.
- The *shuffling* technique uses the existing data and moves the values between rows in such a way that no values are present in their original rows.
- The *number and date variance* technique varies the existing values in a specified range in order to obfuscate them. For example, birth date values could be changed within a range of plus or minus 90 days.
- The [*encryption*](#) technique scrambles the data algorithmically. This technique will not produce realistic-looking data and can make the data larger.
- The *nulling out* technique simply removes the sensitive data by deleting it.
- The *table-to-table synchronization* technique masks data assuring that the results are referentially intact. If two tables contain columns with the same data values and those columns are masked in one table, the second table is updated with the changed value, too.

The bottom line with data masking is that data sanitized by masking is protected from data breaches and hacking. If the data has been masked, it does not matter if thieves gain access to it because it is not useful to them.

Database Archiving for Long-Term Data Retention

Organizations are generating and keeping more data now than at any time in history. Many factors contribute to this reality. One contributing factor is general data growth. According to industry analysts, enterprise data is more than doubling every year. Additionally, as much as 80 percent of that data is not actively used to conduct business.

Why are we producing so much data? Advances in technology have better enabled our ability to capture and store data. But technology alone is not sufficient to account for the current rate of data growth.

Data is retained for both internal and external reasons. Of course, when an organization requires the data to conduct business and make money, that data will be retained. And today's organizations are storing more data for longer periods of time for many internal reasons. Typically, data is stored longer than it used to be in order to enable analytical processes to be conducted on it. Thus, businesses are inclined to keep data around for longer periods of time.

Today's organizations are storing more data for longer periods of time.

But external reasons, typically driven by the mandate to comply with legal and governmental regulations, also compel businesses to store additional data. Indeed, data retention is a significant aspect of regulatory compliance that requires focus and attention. The need to retain data is impacted not just by the normal culprits, such as Sarbanes-Oxley and HIPAA, but also by over 150 international, federal, and local laws that govern how long data must be retained. Organizations need to develop plans for archiving data from the operational databases as their data retention requirements expand over longer and longer periods of time.

The Life Cycle of Data

As data moves throughout its useful life cycle, it progresses through five distinct phases: creation, operational, reference, archived, and discard, as shown in [Figure 15.1](#). Data is created at some point, usually by means of a transaction. For a period of time after creation, the data enters an operational state. The data is required to conduct business. The operational state is followed by the reference state. During this phase data is still needed for reporting and query purposes: internal reports, external statements, or simply

in case a customer asks about it. Then, after some additional time passes, the data is no longer needed for business purposes and it is no longer being queried. But the data must be saved for regulatory and legal purposes. This is the archive state.

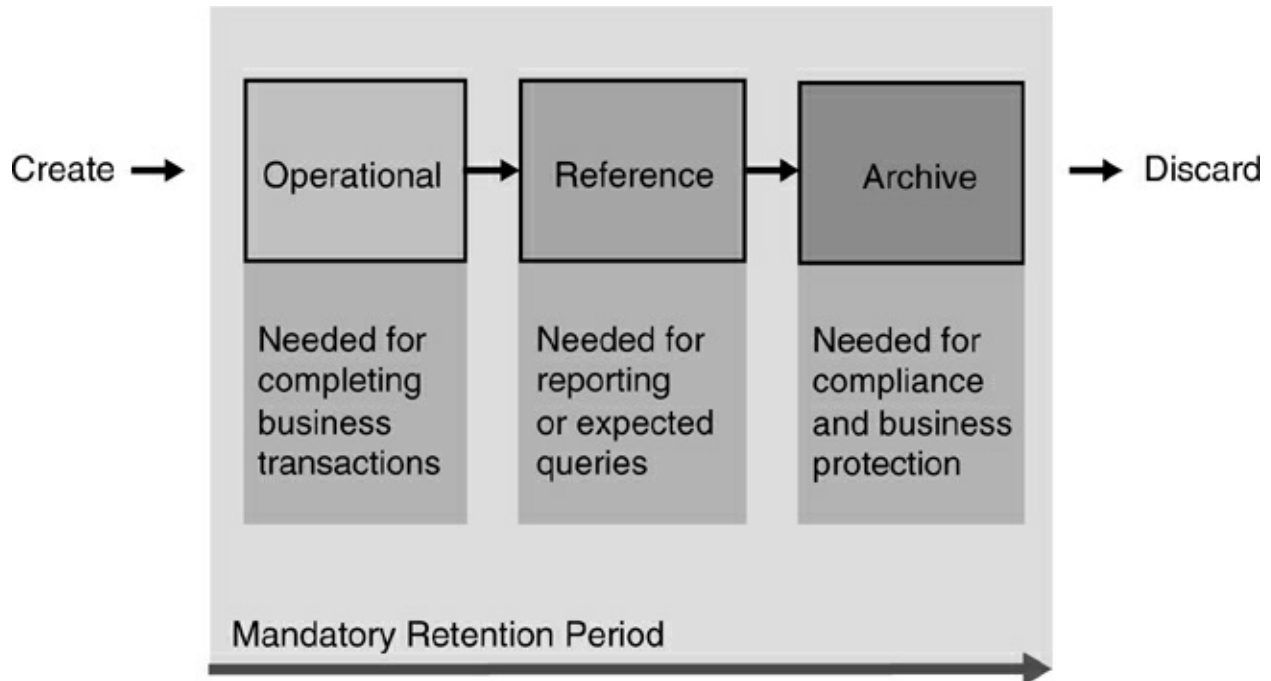


Figure 15.1. The data life cycle

After a designated period of time, the data is no longer needed at all and must be discarded. This phase is often relegated to an afterthought as data continues to pile up but is not used. When there is no legal requirement to maintain data, you should demand that such data be destroyed. If data that is not required for business or regulatory purposes is kept, it becomes a liability instead of an asset. Any data that you keep becomes discoverable and can be used against you in a lawsuit. So why enable anyone to use your data against you if it is not a legal requirement to do so?

The data life cycle states nothing about where the data is stored or what technology is used. But it makes sense to move archive data out of the operational database for many reasons.

Database Archiving

Database archiving is the process of removing selected data records from operational databases that are not expected to be referenced again and storing them in an archive data store where they can be retrieved if needed.

Before moving on, let's break down this definition to make sure it is understood. We say *removing* because the data is deleted from the operational database when it is moved to the archive. If the data is still required for operational requirements, it is not ready to be archived. When data moves into the archive state, query and access are no longer anticipated, so removing it is not problematic.

Next, we say *selected records*. We do not want to archive database data at the file or table level. We need only those specific pieces of data that are no longer needed by the business, but also related data. The archive must be able to selectively choose particular pieces of related data for archival—not the whole database, not an entire table, and not even a specific row. Instead, all of the data that represents a business object is archived at the same time. For example, if we choose to archive order data, we would also want to archive the specifics about each item on that order. This data likely spans multiple constructs within the database.

The next interesting piece of the definition is this: *and storing them [the data records] in an archive data store*. Technicians sometimes confuse data archiving with purging data, but the two are very different (see the sidebar "[Archive versus Purge](#)"). Archived data is stored separately from the operational database and does not require either the DBMS or the applications. Archived data is separate from and independent of the production systems from which it was moved. Archiving confers several benefits upon the database and your environment, including better performance of production systems (because less data needs to be searched) and better protection for archived data (because it is protected in a separate data store).

Archived data should be stored separately from the operational database.

The final component of the definition that warrants clarification is *where they can be retrieved if needed*. The whole purpose of archiving is to maintain the data in case it is required for some purpose. So the data must be readily accessible without requiring a lot of manual intervention.

Archive versus Purge

Archiving data and purging data are two very different processes. Archived

data is removed from the operational data store and maintained in an archived data store. Purged data is removed from the operational data store and discarded.

Do not confuse archiving database data with purging database data. Both can be useful, but they serve completely different purposes.

Determining the Scope of Data Retention

Before you can archive data from your operational databases, you must first determine exactly what must be retained. In other words, data must be archived at the business object level. Another way of saying this is that we need to archive selected records. This is important because not only must those specific pieces of data that are to be archived be retained, but any related data that makes the archived data understandable must accompany it into the archive. The archive needs to be able to selectively choose particular pieces of related data for archival; all of the data that represents a business object is archived at the same time.

For example, if you choose to archive order data, you would also want to archive the specifics about each item on that order. Think about the order entry and management systems you've encountered in the past. You cannot model the database such that all the data is in one table, or at least not if you want to avoid producing junk. There will be an ORDER table and a one-to-many relationship to the ITEM table, because an order can consist of more than one item. If you go to [Amazon.com](https://www.amazon.com) and order three books, there is one order with three items on it, right?

Furthermore, your system will likely track products and customers, so you'll have a CUSTOMER table with a one-to-many relationship going from CUSTOMER to ORDER. This means a customer can place more than one order, but each order is for a single customer only. You'll also have a PRODUCT table with a one-to-many relationship going from PRODUCT to ITEM. This means a product can be on multiple orders, but each individual order will have only one entry for the product. If the customer orders multiples, it will be specified in the quantity field in the item.

Just to make it interesting, let's also include a table called LOCATION that stores inventory information for products, such as bin number and quantity on hand. We'd have a one-to-many relationship going from PRODUCT to LOCATION. This means a product may be stored in more than one bin, but

each bin contains only a single product type. Because a picture is worth a thousand words, refer to [Figure 15.2](#) for an image of this data model.




Figure 15.2. Archive data model

The data spans multiple constructs within the database, meaning tables for DB2 or Oracle and segments and/or databases for IMS. What do we archive? Let's say we need to retain order details for ten years after the order was shipped. Do we really need all of that order data for the full ten years clogging up our operational databases? Probably not. For the purposes of this business, the data is operationally relevant for perhaps a month or two. It might be referenced for reporting or by customers for a year to 18 months after that. So, we decide to keep it in the production databases for two years after the product ships and archive it after that.

But what needs to be archived? If we just archive the ORDER and ITEM tables, we'll be missing pertinent information in the archive. We also need to bring along customer and product information for reference. That is, we have to make sure that the archive contains more than just the CUSTNO in ORDER and the PRODNO in ITEM. We also need the customer name and contact details. What good would CUSTNO 18123546 be if that customer is no longer in the operational database? And would PRODNO 99 be of any use without the remaining product details?

So, we need to set up our archive policy to archive the orders from the ORDER table along with the related items in the ITEM table. This means deleting the data from the operational database when we put it into the

archive, and bringing along related products from PRODUCT and customers from CUSTOMER. This data is moved to the archive with the ORDER and ITEM data, but it is not deleted from the operational database. After all, those products may be in other order items that we are not archiving. Those customers may also have other orders that we are not archiving. So, we would need to create other archive plans for our customer and product data.

A policy is required to define what data should be archived.

Developing accurate and useful database archiving policies is not a trivial matter. Defining what is to be archived and how it is to be archived requires a mix of database skills, business acumen, and knowledge of legal and government regulations. By archiving your data as soon as it is no longer required for operational purposes, you can minimize the risk of data breaches. Of course, this may not be your primary reason to archive data, which is likely to be to preserve your data for long periods of time in a storage construct that is designed for long-term retention and preservation.

Database Archiving Requirements

Let's examine the many capabilities required of a database archiving solution. Perhaps the most important consideration is that the archived data must be *hardware and software independent*. Independence is crucial because of the duration over which the archived data must exist. With a life span of decades (or longer)⁵ it is likely that the production system from which the data was archived will no longer exist—at least not in the same form, and perhaps not at all. Think about the changes that your production applications have undergone over the course of the past ten or 20 years. It is completely unreasonable to expect that today's existing operational environment will exist to enable access to archived data. We constantly change our databases. And the archive must be able to support multiple variations of the data structure as it changes.

The archive solution must be able to store a *large amount of data*. As we store more data, we will archive more data. And when we combine this with long regulatory-mandated data retention periods, we have an explosive combination.

The archive must be able to *manage data for very long time periods*. Many data retention requirements are stated in decades. So the archived data will

outlive the systems and the programmers who generated them. The archive also will outlive the media we store it on. No media lasts forever; consider that the life span of tape is seven years (also, see the sidebar “[Obsolete Hardware](#)”). So, the archive must be able to repurpose the archived data from one type of media to another. And ideally it should do this automatically as the media reaches the end of its useful life.

Data, to support regulatory compliance, must remain unchanged once it is archived. So the archive must be able to *protect against data modification*. Only read access should be available to the archived data (with the exception of periodic administration). Archived data must be guaranteed to be authentic. And mechanisms to prevent surreptitious modification are necessary, too.

To support regulatory compliance, data must remain unchanged once it is archived.

Finally, the archive requires *metadata* to be useful, both metadata defining the archived data as well as metadata defining what to archive and when. The archive must be able to store multiple versions of the first type of metadata. As the operational schema changes, the archive must track and function across these variations in schema. The second type of metadata controls which data is archived, when, and from where. This is the metadata that drives and defines the archive itself. Both types of metadata are needed for the archive to operate.

Obsolete Hardware

When planning for the storage and management of data over long periods of time, be sure to factor in hardware obsolescence as a challenge. It is possible that data may be stored on media that cannot be read by modern hardware.

One consultant tells of a SOX recovery trial run that met such a fate. The company asked its archival vendor to deliver seven-year-old archive tapes. When the staff went to mount the tapes, they discovered that they no longer had that model of tape drive anywhere in the data center. In fact, it was not commercially available anymore. They had to buy several used but working units off of eBay.

Taking all of these considerations into account, then, a secure, durable archive data store must be used to retain data that is no longer needed for operational purposes, and it must enable query retrieval of the archived data in a meaningful format until it is discarded.

Operational databases are no place to maintain historical data over long periods of time. Database archiving will become more prevalent over time, and wise organizations will start planning their database archiving needs today.

Components of a Database Archiving Solution

The diagram in [Figure 15.3](#) depicts the necessary components of a database archiving solution. Starting at the top of the diagram there are two operations that can be performed on databases: data extract and data recall. The extract removes data from the operational database, and the recall restores archived data back to the operational database.

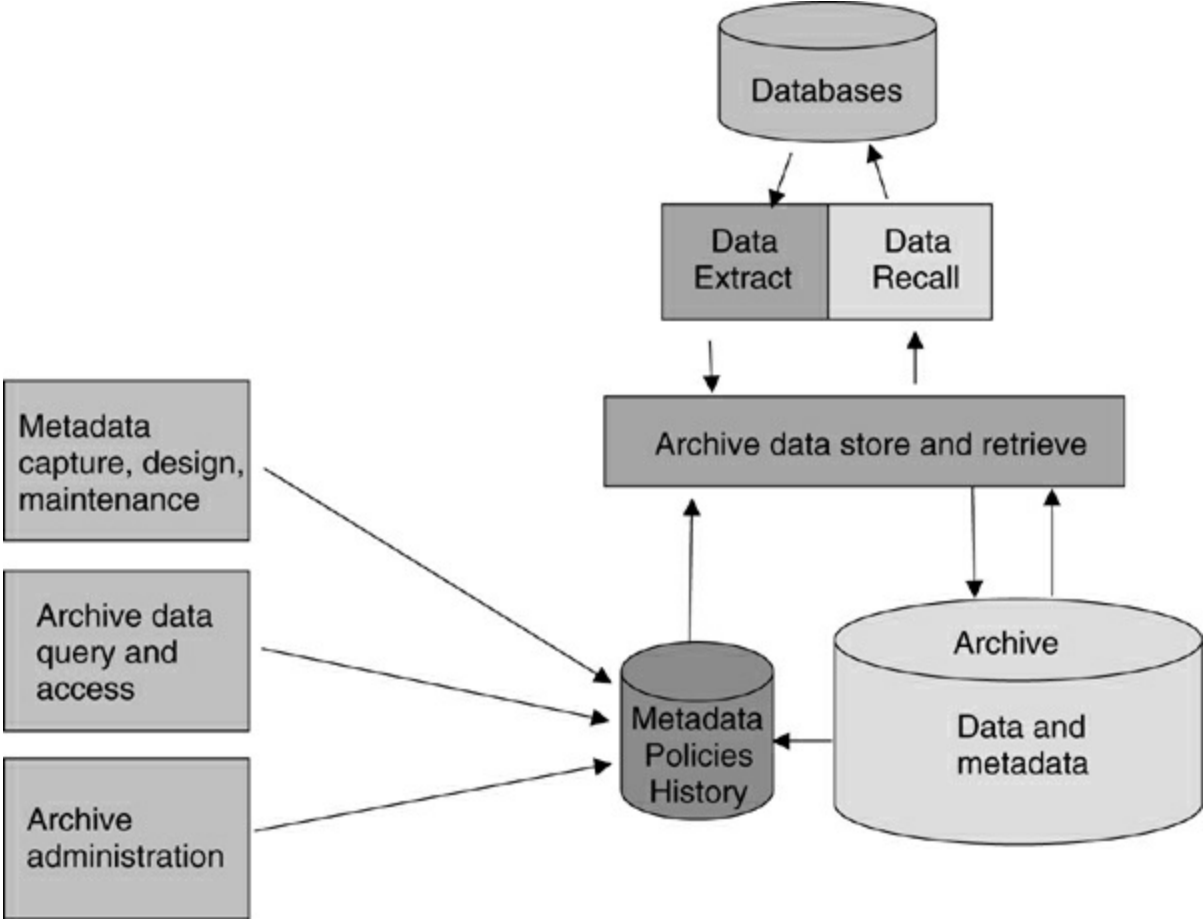


Figure 15.3. Database archival components

The whole process requires metadata to operate. You must capture, validate, and enhance the metadata to drive the archive process. You need to know the structure of the operational database and the structure of the archive. Furthermore, a robust archiving solution will be policy based. This means that policy rules are written to dictate what data is archived, when it is archived, and for how long it must remain in the archive before it is discarded. This policy-based metadata must be maintained and monitored against the archive on an ongoing basis.

A query capability allowing direct reads against the archive is important, too. Queries against an archive data store will not necessarily be the most efficient access because of differences in the metadata over time. However, queries against archived data typically are not very performance sensitive, at least not like typical transaction data.

Finally, it is also important to have an ongoing maintenance capability for the archived data. This encompasses administrative tasks such as security, access audit, administration of the structures (such as reorganization), backup and recovery, and so on.

The Impact of e-Discovery on DBA

One of the looming issues facing data management professionals is preparation for e-discovery. Although regulations mandate that we retain data longer, there are rules and regulations that dictate when and how organizations will need to access and produce data that is retained. The only reason that data is being kept is because there may be a need to see it again.

The ability to produce retained data upon request is typically driven by lawsuits. You probably can recall examples of courtroom showdowns on television where truckloads of paper documents were required during the discovery process of the lawsuit. But times have changed. Increasingly, the data required during the discovery process is electronic, not written. That is, the data is stored on a computer, and much of that data is stored in a database management system.

Increasingly, the data required during the discovery process is electronic.

Which brings me to the Federal Rules of Civil Procedure (FRCP), which are the rules used by U.S. district courts to govern legal proceedings. One of

the items in this set of rules dictates policies governing discovery. Discovery is the phase of a lawsuit before the trial occurs during which each party can request documents and other evidence from other parties or can compel the production of evidence.

The FRCP underwent changes in late 2006 to focus more on electronic data. Rule 34b of the FRCP was changed to state that “A party who produces documents for inspection shall produce them . . . as they are kept in the usual course of business. . . .” This change clearly compels organizations to improve their ability to produce electronic data.

And Rule 37 of the FRCP adds a new section, Rule 37(f), which provides a safe harbor from sanctions arising from spoliation. According to this section, “Absent exceptional circumstances, a court may not impose sanctions under these rules on a party for failing to provide electronically stored information as a result of the routine, good faith operation of an electronic information system.” Basically, this section shines a spotlight on the need for organizations to develop a clearly articulated, well-executed, and uniformly enforced records retention program. And that program should include database data. Instituting policies and procedures for how data is treated for long-term retention can provide some level of protection from “adverse inference” rulings arising from spoliation.

There are likely to be additional implications of the changes to the FRCP, especially when coupled with the growing list of data breaches and the growing regulations being voted into law by federal and state governments. It means that we will be forced to treat data as the corporate asset that it is—instead of just saying that we treat it that way.

Closer Tracking of Traditional DBA Tasks

Regulatory compliance also can help to improve the focus on traditional DBA procedures and processes. Of course, DBAs managed corporate database systems prior to the recent increase in governmental and industry regulations. But the accounting and regulatory demands of the regulations often require those DBA practices to be strengthened and formalized. Specifically, database security and access control practices may need to be modified. And a more comprehensive approach to the auditing of data access may be required, especially as it relates to unusual or suspect activity. But this also is true for tried-and-true DBA practices such as change management

and database backup and recovery.

Database Change Management

Many organizations do not adhere to strict, formalized change policies for database objects. Furthermore, logging the series of commands needed to implement (or undo) changes may need to be enhanced, if such logging is implemented at all.

Many regulations require all changes to be logged, specifying who did what and why each change occurred. Too often, DBAs simply issue ALTER commands to change the database as needed in reaction to performance issues or development requirements. Although a robust database auditing solution can capture such changes, it would be a much better practice to require database changes to be made using a change management tool that automatically tracks changes.

A best practice is to require database changes to be made using a change management tool that automatically tracks changes.

DBAs have had to deal with change management as long as databases have been used. But in many cases these tasks were attacked in a low-cost, ad hoc manner. Maybe there was not sufficient capital to expend on DBA processes; maybe the DBAs were capable enough to change database structures without the aid of tools. In many cases, with the stricter rules of compliance in place, this is no longer acceptable.

Database Backup and Recovery

When data professionals think about regulatory compliance, we tend to consider first the data in our production databases. After all, it is this data that runs our business and that must be protected. So we work to implement database auditing to know who did what to which data when; or we tackle database security and data protection initiatives to protect our data from prying eyes; or we focus on improving data quality to ensure the accuracy of our processes. These are all worthwhile endeavors, but focusing exclusively on active production data is insufficient to ensure compliance. The recoverability of the database is a compliance issue, too. Improved backup and recovery practices and procedures must be an essential component of your compliance plans.

Ensuring the integrity and availability of your databases is the primary focus of backup and recovery planning. But what about database recovery in terms of compliance and regulations? Let's examine database recovery through the lens of COBIT (see the sidebar titled "[What Is COBIT?](#)").

What Is COBIT?

COBIT is a framework of IT best practices that companies can use to improve management of their IT organizations, to improve the value of IT, and to ensure that the goals of the IT organization are aligned with the goals of the business. The current COBIT guidelines are available from the IT Governance Institute at www.isaca.org/Knowledge-Center/cobit/Pages/FAQ.aspx. The IT Governance Institute was created by ISACA, the Information Systems Audit and Control Association.

COBIT is about recognizing and safeguarding the value of information as a corporate asset by identifying and managing risks and ensuring corporate governance via effective controls. The crux of COBIT is to link IT and business goals, to identify responsibilities of business and IT owners, and to monitor performance, evaluating it against metrics and maturity models.

The COBIT framework consists of 34 specific control objectives, organized into four domains: Plan and Organize (PO), Acquire and Implement (AI), Deliver and Support (DS), and Monitor and Evaluate (ME). The COBIT practices are business focused, process oriented, measurement driven, and control based. Best practice frameworks like COBIT are vital tools for ensuring compliance with regulations such as Sarbanes-Oxley.

Do not confuse COBIT with frameworks such as ITIL. ITIL is an IT Service Management framework for identifying, planning, delivering, and supporting IT services to the business. Think of COBIT as describing "what IT needs to do" and ITIL as describing "how IT can deliver" those services.

COBIT and Recovery

Database recovery needs to be tackled from a best practice approach to enable your organization to do the kind of up-front planning and routine

monitoring and evaluation that COBIT advocates. An organization that has adopted COBIT as a best practice framework understands the critical value of information to the business, and the need to assure its integrity and availability.

Of course, compliance is not the overarching factor compelling DBAs to develop backup policies and procedures for all database objects. This should be driven by business availability requirements, and most DBAs have done this, at least to some degree. But what most have not done is implement regular systematic checks for the ongoing viability of their backup and recovery plans to match their recovery time objectives—or even to ensure that their existing backups are valid and could be used in a recovery situation.

Recoverability is addressed by the following 19 COBIT objectives across three process domains:

- PO9.4: *Assess risks*. During planning and organization you must assess the risk of databases being unrecoverable from backups.
- DS1.3 and DS1.4: *Define and manage service levels*. Metrics are required to define service-level objectives for recovery. Do you know how long it would take to recover a specific database object (or series of objects)? If not, how can you assure that application service levels will be met or exceeded?
- DS3.2, DS3.3, DS3.4, DS3.5, and DS3.8: *Manage performance and capacity*. Regularly checking the health of your recovery aids capacity management by improving the availability of information and the IT resources that depend on it.
- DS 4.10, DS4.11, and DS4.12: *Ensure continuous service*. Again, ensuring service is impossible without being able to ensure recoverability (including data that is mirrored to backup IT sites and/or off-site backup data stores).
- DS11.9, DS11.19, DS11.20, DS11.21, DS11.23, and DS11.24: *Manage data*. Any number of issues may require recoverability as part of an ongoing data management effort. COBIT Objective DS11.24 specifically covers verifying the usability of backups.
- M1.1 and M1.2: *Monitor the processes*. Ongoing monitoring of recoverability is needed to verify every backup job and its effectiveness in your environment (logging, memory, system resources, etc.).

Organizations need to acquire tools and implement procedures that help to verify the integrity of their backups, the system settings that could affect their ability to recover, and the processes associated with backup and recovery of their databases. Analyzing your database system, data, and backups and determining their health and usability should be a regular practice. If not undertaken, a system failure, logical error, malicious destruction, or catastrophic event could render your databases unusable, impact your business, and maybe even threaten the ongoing viability of your business.

Organizations need tools and procedures that help to verify the integrity of database backups.

Summary

Regulatory compliance places renewed emphasis on many traditional DBA tasks while also expanding the meaning of some of them. Backup and recovery, metadata management, and change management require additional focus, whereas features like database auditing and database archiving may be required for the first time. This is causing additional workload for database administration.

Fortunately, though, now that high-level corporate executives have to vouch for the accuracy of company data, tools that can help to assure data accuracy and integrity are no longer a luxury but are required to avoid prosecution.

Review

1. Ensuring compliance requires a collaborative effort among which three groups?
2. Name two traditional DBA duties that may require additional attention in the light of a regulatory compliance project.
3. Prosecution for compliance failure is not just an idle threat: true or false?
4. Define *database archiving*. Contrast it with *data purging*.
5. Name the different techniques that can be used to implement database auditing.
6. Why are metadata and data quality important aspects of regulatory

- compliance with respect to database systems and applications?
7. What is the purpose of data masking?
 8. What is a potential drawback of implementing a database auditing solution?
 9. What is PII?
 10. Some organizations have identified the need to retain data for more than 100 years: true or false?

Suggested Reading

- Afyouni, Hassan A. *Database Security and Auditing: Protecting Data Integrity and Accessibility*. Boston, MA: Thomson (2006). ISBN 0-619-21559-3.
- Anand, Sanjay. *Sarbanes-Oxley Guide for Finance and Information Technology Professionals*. 2nd ed. Hoboken, NJ: John Wiley & Sons (2006). ISBN 0-471-78533-9.
- Ben Natan, Ron. *Implementing Database Security and Auditing*. Oxford, UK: Elsevier Digital Press (2005). ISBN 978-1-55558-334-7.
- Cougias, Dorian J., et al. *Say What You Do: Building a Framework of IT Controls, Policies, Standards, and Procedures*. Lecanto, FL: Schaser-Vartan (2007). ISBN 978-0-9729039-6-7.
- Denning, Dorothy E. *Cryptography and Data Security*. Reading, MA: Addison-Wesley (1983). ISBN 0-201-10150-5.
- Fowler, Kevvie. *SQL Server Forensic Analysis*. Boston, MA: Addison-Wesley (2009). ISBN 978-0-321-54436-0.
- IT Governance Institute. *CobiT 4.0: Control Objectives, Management Guidelines, Maturity Models*. Rolling Meadows, IL: IT Governance Institute (2005). ISBN 1-933284-37-4.
- Mack, Mary, and Matt Deniston. *A Process of Illumination: The Practical Guide to Electronic Discovery*. Portland, OR: Discovery Center of Excellence (2004). ISBN 0-9725542-1-1.
- Olson, Jack E. *Database Archiving: How to Keep Lots of Data for a Very Long Time*. Burlington, MA: Morgan Kaufmann (2009). ISBN 978-0-12-374720-4.

- . *Data Quality: The Accuracy Dimension*. San Francisco, CA: Morgan Kaufmann (2003). ISBN 1-55860-891-5.
- Ottman, John B., Jr. *Save the Database, Save the World!* New York, NY: Sumo Press (2010). ISBN 978-1-4583-6368-8.
- Perry, William E. *Control in a Data Base Environment*, Wellesley, MA: QED Information Sciences (1980). ISBN 0-89435-042-0.
- Redman, Thomas C. *Data Quality: Management and Technology*. New York, NY: Bantam (1992). ISBN 0-553-09149-2.
- . *Data Quality: The Field Guide*. Boston, MA: Digital Press (2001). ISBN 1-55558-251-6.
- Reynolds, George W. *Ethics in Information Technology*. 3rd ed. Boston, MA: Course Technology (2010). ISBN 978-0-538-74622-9.
- Schuler, Karen, et al. *E-discovery: Creating and Managing an Enterprisewide Program*. Burlington, MA: Syngress (2009). ISBN 978-1-59749-296-6.
- Weill, Peter, and Jeanne W. Ross. *IT Governance: How Top Performers Manage IT Decision Rights for Superior Results*. Boston, MA: Harvard Business School Press (2004). ISBN 1-59139-253-5.
- Wright, Craig, et al. *The IT Regulatory and Standards Compliance Handbook: How to Survive an Information Systems Audit and Assessments*. Burlington, MA: Syngress (2008). ISBN 978-1-59749-266-9.

16. Database Backup and Recovery

Things break. When a new database is created or a new application goes online, everything is fresh and new—and running as designed. But the environment changes over time. New hardware and software are added, along with more users, more data, more requirements—more, more, more. Systems wear down as they are taxed on a daily basis to deliver service.

The DBA must be prepared for situations where a failure impacts the availability, integrity, or usability of a database. Reacting to failures and service disruptions is a key component of the DBA's job. The ability of the DBA to react accordingly depends directly on having a well-planned approach to database backup and recovery.

Reacting to failures and service disruptions is a key component of the DBA's job.

The Importance of Backup and Recovery

I know that many DBAs believe that ensuring optimal database and application performance is the most important task in their list of job responsibilities, but it is not true. These DBAs are confusing frequency with importance.

Most DBAs probably are conducting performance-related tasks more often than building backup plans—and they had better be managing performance more frequently than they are actually recovering their databases or their company has big problems! But recoverability should be at (or near) the very top of the DBA task list, definitely before performance. Why? Well, if you cannot recover your databases after a problem, it won't matter how fast you can access them, will it? Anybody can deliver fast access to the wrong information. It is the job of the DBA to keep the information in the company's databases accurate, secure, and accessible.

Anybody can deliver fast access to the wrong information.

So what do DBAs need to do to assure the availability and accuracy of database data? This chapter will outline the basics of database backup and

recovery and provide guidance on building a robust backup and recovery plan for your databases.

Preparing for Problems

Numerous daily hazards can cause system failures. As you plan your database backup and recovery strategy, be sure to consider all of these various threats to database integrity and availability. Of course, it is wise to take precautionary measures to prevent failures. Techniques such as UPS systems, mirrored disks, and failover technology can minimize the need to recover, but no amount of planning and regulation can prevent unexpected failures.

Database failures that may require recovery can be divided into three categories:

- *Instance failures* are the result of an internal exception within the DBMS, an operating system failure, or other software-related database failure. In some cases, an instance failure can result in corruption of data that requires a recovery, but usually such failures do not damage data, so the DBMS simply needs to be restarted to reestablish normal operations.
- *Application (or transaction) failures* occur when programs or scripts are run at the wrong time, using the wrong input, or in the wrong order. An application failure usually results in corrupt data that requires a database restore or recovery. The sooner an application failure is identified and corrected, the smaller the amount of damage to the database will be.
- *Media failure* is likely to damage data, too. Media failure includes damage to disk storage devices, file system failures, tape degradation or damage, and deleted data files. Although less common in practice, damaged memory chips also can cause data corruption. After a media failure, the database will likely be in a state where valid data is unreadable, invalid data is readable, or referential integrity is violated. Outages due to media failures can often be avoided by implementing modern disk technologies such as RAID, which is covered in more detail in [Chapter 18](#), “[Data and Storage Management](#).”

It is common for organizations to manage a terabyte or more

of data on a single database server.

Businesses today are demanding higher throughput and around-the-clock availability even as they increase the amount of stored and processed data. Not too long ago, we talked about data in terms of gigabytes. Now it is common for organizations to manage a terabyte or more of data on a single database server. Therefore, more data needs to be constantly available and has to be processed faster than ever before. Organizations rely on data to conduct business, so it is imperative that you be prepared with a plan to counteract failures. A sound backup and recovery plan can be thought of as an insurance policy for your data.

Backup

A fundamental component of a database backup and recovery plan is creating backup copies of data. When an error occurs that damages the integrity of the database, a backup copy of the data can be used as the basis to recover or restore the database. However, the full story on backing up a database is not quite that simple.

Backing up databases involves making consistent copies of your data, usually in the form of image copies, which are the output of a COPY utility. The name of the copy utility will vary from DBMS to DBMS. Common names for the backup utility include BACKUP, COPY, DUMP, and EXPORT. Some DBMSs rely on the native operating system's file system commands for backing up data. However, even if the DBMS supplies an internal backup option, the DBA may choose to use facilities that operate outside the realm of the DBMS.

Backing up databases involves making consistent copies of your data.

Current and accurate image copies provide the foundation for database recovery. The DBA must assure the currency and accuracy of the image copies and base the backup plan on the recovery needs of the applications. The DBA will use those recovery requirements to determine how often to take image copy backups and how many backup generations must be kept on hand. Furthermore, the DBA must make sure that the appropriate log records are available or backed up for recovery purposes. To decide the frequency

with which to back up a database object, consider how much time will be needed to recover that object. Duration of recovery is determined by factors such as

- The number of log records that must be processed to recover
- Whether the log is compacted or compressed
- Whether the image copy backup is encrypted or compressed
- The time it takes an operator to mount and dismount the required tapes
- The time it takes to read the part of the log needed for recovery
- The time needed to reprocess changed pages

Additionally, recovery duration depends on the architecture of the DBMS. For example, mainframe DB2 keeps track of log range information and reads only the required log files for any recovery operation. However, some DBMSs require that all the log files be read to scan for information needed for recovery.

Keep in mind that database backups taken while there is read-only activity, or no activity, can be restored back to that point in time using only the backup—no log files are required. This can simplify and minimize the cost of a recovery.

In general, the more often you make an image copy, the less time recovery takes. However, the amount of time required to make an image copy backup must be balanced against the need for concurrent processing during the backup process.

The DBA must decide how many complete generations of backups to keep.

The DBA must decide how many complete generations of backups (for both database object copies and log copies) to keep. By keeping extra generations, you can sometimes recover from an error during recovery by switching to an older backup. At a minimum, the retention period should be at least two full cycles. When you make a scheduled full image copy of a tablespace, the latest full image copy and at least one other previous copy will exist, plus the log created since the earliest image copy. You may want to consider keeping at least two weeks' worth of copies. Of course, the number of copies you decide to keep must be tempered by the number of associated

logs that must also be maintained for the backups to remain viable.

The following guidelines on making image copy backups will help assure a recoverable environment:

- Make at least two local copies of each image copy backup to help avoid an unrecoverable state in the case of a media error (for example, a damaged tape).
- Coordinate your local backup strategy with your disaster recovery backup strategy. Many backup utilities permit both local and off-site backups to be created simultaneously.
- Keep at least two generations of image copy backups for each database object. If the most recent image copy fails, you can fall back to the older copy and still be able to recover.
- Consider creating image copy backups to disk, and then migrating them to tape (or optical disc, such as CD or DVD), which can speed up the image copy process. Not only is disk faster than tape, but also the recovery will not have to wait for a manual tape mount.
- When image copy backups are migrated to tape, consider compressing the files to reduce the number of tapes needed for large image copy backup files. This usually can be accomplished using the facilities of the tape drive. Also, consider encrypting the backup files for sensitive data.
- Be sure to include the system catalog database objects in your backup and recovery plans. System catalog data should be backed up at least weekly—perhaps more frequently for very dynamic systems. Whenever you issue DDL, it changes the data in the system catalog. Failing to back up the system catalog after a DDL change may cause the changed database objects to become inaccessible after a recovery.
- Ensure that the backup process is restartable. For example, consider a database backup process that takes three hours to complete. If the process fails after two and a half hours, the restart should require only a half hour to finish. If the backup were not restartable, it would have to start again from the beginning.
- After the backup has completed, use the DBMS's facilities to verify the correctness of the backup, for example, the DB2 db2ckbkp operation or the Sybase BCP utility.

- Data that is not stored in a database but is used by database applications should be backed up at the same time as the database objects.

As a rule of thumb, make image copy backups infrequently enough that you do not interrupt daily business, but frequently enough that you can afford the time required to recover the database object. Of course, this is easy to say but hard to accomplish. You will need to keep abreast of the application requirements of each database object in order to strike the right balance for creating image copy backups.

Certain database activities make it a wise decision, if not a requirement, to take a full image copy backup of a database object. For example, before reorganizing the system catalog, you should make image copy backups of all system catalog database objects.

It is often a wise decision to take a full image copy backup of a database object.

In addition, after every reorganization you should take a full image copy backup of the reorganized database object. If you need to recover the object, you can recover to the reorganized version of the database objects instead of a prior, disorganized version. Alternatively, if the reorganization was logged, you will not need to apply all the log records written during the reorganization.

Another situation requiring an image copy backup is when data is loaded into a table with logging disabled. After you have loaded data into a table using the load utility, immediately make a full image copy backup. Even if you load with logging enabled, it is wise to create a full image copy backup of the data after the load process completes. That way, if you need to recover the object, you will not need to apply the log records written during the load.

Finally, image copy backups should be taken after point-in-time recoveries. This will ensure that the database object can be recovered if another failure occurs.

Full versus Incremental Backups

Two types of image copy backups can be taken: full and incremental. As a DBA, you will need to learn the difference between the two and implement the proper image copy backup strategy based on application needs and database activity.

There are two types of image copy backups that can be taken: full and incremental.

A *full image copy backup* is a complete copy of all the data in the database object at the time the image copy was run. An *incremental image copy backup*, sometimes referred to as a *differential backup*, contains only the data that has changed since the last full or incremental image copy was made. The advantage of taking an incremental backup rather than a full backup is that it can sometimes be made more quickly, and it occupies less space on disk or tape. The disadvantage is that recovery based on incremental copies can take longer because, in some cases, the same row is updated several times before the last changes are restored.

For example, suppose you took a full image copy of a database object early Monday morning at 2:00 a.m. and then took an incremental image copy at the same time the following three mornings (see [Figure 16.1](#)). The full image copy plus all three incremental image copies need to be applied to recover the tablespace. If the same column of the same row was updated on Tuesday to “A,” Wednesday to “B,” and Thursday to “C,” the recovery process would have to apply these three changes before arriving at the final, accurate data. If a full image copy were taken each night, the recovery process would need to apply only the latest image copy backup, which would contain the correct value.

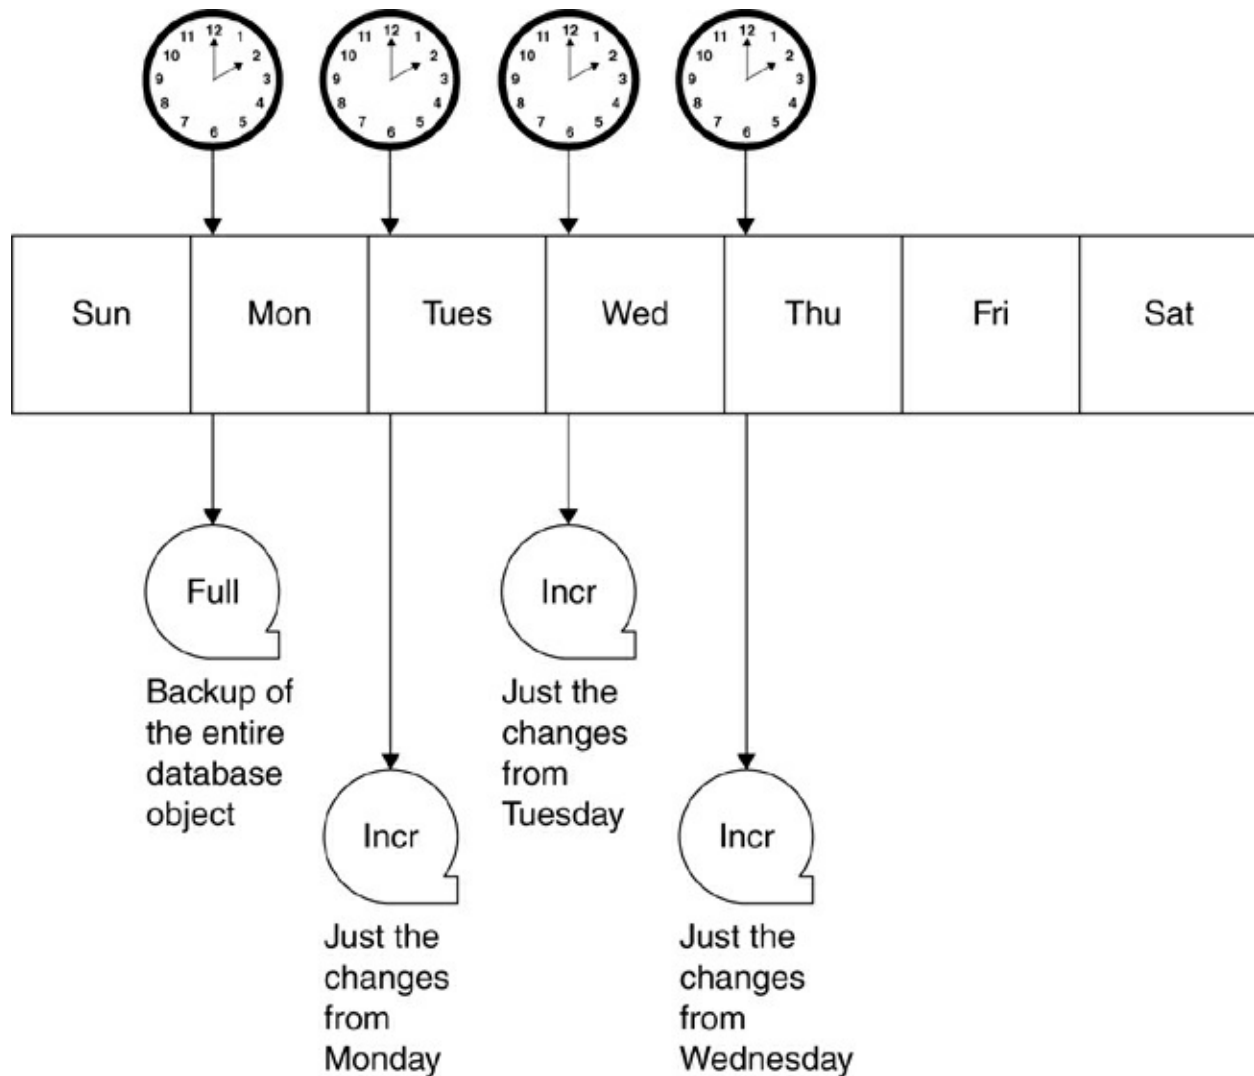


Figure 16.1. Image copy backups

Some DBMSs provide the capability to analyze a database object to determine if a full or incremental backup is recommended or required. This is typically accomplished using an option of the copy utility. If such an option exists, the DBA can run the copy utility to examine the amount of data that has changed since the last image copy backup was taken. Furthermore, the DBA can set a threshold such that a full image copy is taken when more than a specified amount of data has changed; an incremental image copy is taken when the amount of data that has changed is less than the threshold. When this option is not available, the DBA will need to set up the type of image copy backups to be taken based on his knowledge of the applications and their use of the databases.

Most database objects will be set up for full image copy backups instead of incremental image copies. The time saved by incremental copying is useful only when a small amount of data has been changed. However, the additional work required for a recovery using incremental copies can be troublesome. In general, full image copy backups are recommended for database objects where 30 to 40 percent or more of the data blocks are modified between backups. The DBA makes this determination based not just on data volatility but also on factors such as criticality of the data, availability requirements, and the functionality of the DBMS.

Favor full image copies for small database objects.

Favor full image copies for small database objects. The definition of “small” will vary from site to site and DBMS to DBMS. For large, mainframe-based databases, a “small” database object is probably within the 150GB to 200GB range. For smaller, though still significant UNIX-based databases, a “small” database object is probably about 100GB to 150GB. For small workgroup or Windows-based database sites, even a 100GB database object is on the large side.

Consider using incremental image copies to reduce the batch processing window for very large database objects that are minimally modified between image copy backups. The DBA must base the full-versus-incremental decision on the percentage of blocks of data that have been modified, not on the number of rows that have been modified. In general, consider incremental copying as the tablespace becomes larger and the batch window becomes smaller.

Additionally, some scenarios are not compatible with incremental image copy backups. Some DBMSs permit the user to disable logging during some operations and utilities. Whenever an action is taken that adds or changes data without logging, a full image copy is required.

Some scenarios are not compatible with incremental image copy backups.

Merging Incremental Copies

If the DBMS supports incremental image copy backups, it may also support incremental copy merging. A merge utility, sometimes referred to as

MERGECOPY, can be used to combine multiple incremental image copy backups into a single incremental copy backup, or to combine a full image copy backup with one or more incremental image copy backups, to create a new full backup.

If your DBMS supports merging incremental copies, consider running the merge utility to create a new full image copy directly after the creation of an incremental copy. The merge utility can be run with no impact on concurrent data access or the batch window. If you wait until recovery is required to run the merge, downtime will be increased because the merge (or similar processing) will occur during the recovery process while the database object is unavailable.

Database Objects and Backups

Typically, an image copy backup is made at the database, tablespace, or table level. The level(s) supported will depend on the DBMS being used. In general, though, the idea is to back up the database object or objects that contain the data. In general, the more granular control the DBMS provides for backup of database objects, the easier it will be to effectively implement a useful backup and recovery strategy.

An image copy backup is made at the database, tablespace, or table level.

Copying Indexes

Some DBMSs support making backup copies of indexes. Indeed, some DBMSs require indexes to be backed up, whereas index backup is optional for others. Index backup can be optional because the DBMS can rebuild an index from the table data. Therefore, a typical recovery scenario would involve recovering tables or tablespaces and then using the table data to rebuild the indexes. If the DBMS supports index backups, you can choose to copy indexes and then recover using the image copy backups of the indexes.

You will need to examine the trade-offs of copying indexes.

As a DBA, though, you will need to examine the trade-offs of copying indexes if your DBMS supports index backup. The question DBAs must answer for each index is “Rebuild or recover?” The more data that must be indexed, the longer an index rebuild will require in a recovery situation. For

larger tables, backing up the index can result in a much quicker recovery—although at the expense of the increased time required for backup. When multiple indexes exist on the large table, backing them up, again, leads to faster recovery. However, keep in mind that index backups will require additional time to execute during your regular backup process. As a DBA, you will need to weigh the cost of recovery against the cost of backup in making your decision.

Be sure to perform data and index backups at the same time if you choose to back up rather than rebuild your indexes. In other words, when you back up a table (or tablespace or database), be sure to also back up all associated indexes. Failure to do so can result in indexes that do not match the recovered data—which will cause applications to be unavailable or, worse, to receive invalid results.

DBMS Control

The degree of control the DBMS asserts over the backup and recovery process differs from DBMS to DBMS. Some DBMSs record backup and recovery information in the system catalog. That information is then used by the recovery process to determine the logs, log backups, and database backups required for a successful recovery. DB2 for z/OS works this way; refer to the sidebar “[The DB2 COPY Utility](#).” The more information the DBMS maintains about image copy backups, the more the DBMS can control proper usage during recovery.

The DB2 COPY Utility

The COPY utility is used by DB2 for z/OS to create image copy backups. This utility maintains a catalog of image copy information in the system catalog. Every successful execution of the COPY utility causes DB2 to record information in the system catalog that indicates the status of the image copy, the image copy data set name and file details, the date and time of the backup, and log information. This information is read by the DB2 RECOVER utility to enable automated tablespace and index recovery. Only valid image copies recorded in the system catalog can be used by DB2 for recovery.

As time passes, image copy backups become obsolete. New backup copies are made and database objects are recovered to various points in

time. The DB2 DBA must maintain the information in the system catalog because outdated and unnecessary backup rows in the system catalog can slow down the recovery process. Backup information in the system catalog is removed by the DBA using the DB2 MODIFY utility.

On the other hand, some DBMSs do not record backup and recovery information in the system catalog. If this is the case, the DBA must track image copy backup files and assure their proper usage during a recovery. Microsoft SQL Server is an example of a DBMS that works this way. Whenever a backup is made, the DBA must trap pertinent details of the backup such as the type, the time, and the date. During a recovery situation, the DBA must supply the backups (for both database objects and logs) in the proper sequence to assure a valid recovery.

Some DBMSs do not record backup and recovery information in the system catalog.

Backup and recovery for Oracle databases can be handled this way, too; however, most organizations will not choose to do so. Oracle provides a comprehensive method for managing backup and recovery called RMAN. RMAN, which stands for Recovery Manager, is a utility that establishes a connection with a server session and manages the data movement for backup and recovery operations. When using RMAN, Oracle backup and recovery operates similarly to DB2 for z/OS. Refer to the sidebar “[Using Oracle RMAN for Backup and Recovery](#)” for more details.

Concurrent Access Issues

Some backup techniques permit concurrent write access to the database object. Using these techniques allows you to keep the data online during the backup process, but it will slow down any subsequent recovery because the DBMS has to examine the database log to ensure accurate recovery.

Using Oracle RMAN for Backup and Recovery

RMAN is a powerful program for managing the backup and recovery of Oracle data. The DBA can use RMAN to specify files or archived logs to be backed up using the RMAN BACKUP command. Doing so causes RMAN to create a backup set as output. A backup set is one or more data

files, control files, or archived redo logs that are written by RMAN in proprietary format. The only way to recover using the backup set is to use the RMAN RESTORE command. Of course, the DBA can choose to use the COPY command instead. This creates an image copy of a file that is usable outside the scope of RMAN.

RMAN accesses backup and recovery information from either the control file or the optional recovery catalog. The recovery catalog is similar to the DBMS system catalog, but it contains only backup and recovery metadata.

Keep in mind that RMAN is a vehicle for controlling the backup and recovery of Oracle data. It does not, in and of itself, perform the backup and recovery; instead, it controls when and how backups are taken and can be used to coordinate recovery procedures and tasks.

RMAN generally is preferable to other Oracle backup and recovery methods because it is easier to use and provides more functionality. For example, RMAN provides the ability to create incremental backups. Only full image copy backups are available when using traditional Oracle backup and recovery methods.

Some DBMSs and recovery utilities provide the capability to perform change accumulation. The change accumulation process creates an up-to-date image copy backup by merging existing image copies with data from the database logs. This is similar to the merging of incremental image copies.

Some image copy backup techniques allow only read access to the database object. Before the image copy is taken, a QUIESCE (explained in the next section) of the database object establishes a point of consistency. Backups that allow only read access provide faster recovery than those that allow concurrent read-write because the database log is not needed to ensure a proper recovery.

Finally, some image copy backup techniques require the database object to be stopped, or completely offline. Under such conditions, no one can read or write to the tablespace. This type of copy provides fast backup because there is no contention for the tablespace.

The DBA must understand the backup capabilities of each DBMS in the organization.

The DBA must understand the backup capabilities of each DBMS in the organization and plan a proper backup strategy that takes into consideration

- The need for concurrent access and modification during the backup process
- The amount of time available for the backup process and the impact of concurrent access on the speed of backing up data
- The speed of the recovery utilities
- The need for access to the database logs

Some DBMSs use the terms *hot backup* and *cold backup* to describe the concurrent access that can occur while the data is being backed up. A cold backup is accomplished by shutting down the database instance and backing up the relevant database files. A hot backup is performed while the database instance remains online, meaning that concurrent access is possible. Depending on the capabilities of the DBMS you are using, hot backups can be problematic because

- They can be more complex to implement
- They can cause additional overhead in the form of higher CPU, additional I/O, and additional database log archivals
- They can require the DBA to create site-specific scripts to perform the hot backup
- They require extensive testing to ensure that the backups are viable for recovery

Backup Consistency

Be sure your backup plan creates a consistent recovery point for the database object. In order to ensure backup consistency, you need to be aware of all relationships between the database objects being backed up and other database objects. These include application-enforced relationships, referential constraints, and triggers. If you use an image copy backup to recover a database object to a previous point in time, you will need to recover any related database objects to the same point in time. Failure to do so will most likely result in inconsistent data.

If your DBMS provides a QUIESCE utility, use it to establish a point of consistency for all related database objects prior to backing them up. The QUIESCE utility will halt modification requests to the database objects to

ensure consistency and record the point of consistency on the database log. Use the QUIESCE utility even when some database objects do not need to be copied, so that you can provide a consistent point of recovery for all referentially tied tablespaces.

If the DBMS does not provide a QUIESCE option, you will need to take other steps to ensure a consistent point for recovery. For example, you can place the database objects into a read-only mode, take the database objects offline, or halt application processes—at least those application processes that update the related database objects.¹

When to Create a Point of Consistency

If possible, the DBA should create a point of consistency during daily processing. A point of consistency can come in handy if a point-in-time recovery is required. You should consider creating a point of consistency in the following situations:

The DBA should create a point of consistency during daily processing.

- *Before archiving the active log.* If you ever lose your active logs, and you need to use your archive logs for recovery, you can safely apply log records only up to the last recovery point. If you apply logs after that point, you can end up with inconsistent data. If the active logs are not automatically archived, you can quiesce your objects just before you off-load the active log, which will allow you to use the data all the way to the end of the log without being left in an inconsistent state.
- *Before copying related database objects.* Consider creating a point of consistency before copying sets of related tables. This assures that the image copy backups for all of the related database objects are consistent with each other.
- *Just after creating an image copy backup.* When making image copy backups concurrently with online database changes, you can create a point of consistency after the backup is taken to establish a good recovery point.
- *Just before heavy database modification.* When batch jobs are run out of order, or online programs work incorrectly, partial recoveries are used to restore the data to its state preceding the update activity. By

establishing a point of consistency before the update activity, you can consistently recover to that point without having to make image copy backups prior to the heavy update activity.

- *During quiet times.* Establishing a point of consistency can be disruptive during periods of heavy activity. The DBMS must establish a point in time when none of the database objects in question are being updated. By scheduling a QUIESCE during quiet periods, you can avoid such disruptions.

Log Archiving and Backup

All database changes are logged by the DBMS to a log file commonly called the *transaction log* or *database log*. Log records are written for every SQL INSERT, UPDATE, and DELETE statement that is successfully executed and committed. It is possible to review and either recreate or undo the effects of every change made to the database by using the database log.

As time progresses and the number of database changes grows, the database log will increase in size. The database log to which records are currently being written is referred to as the *active log*. When the active database log is filled, the DBMS invokes a process known as log archival or log off-loading. When a database log is archived, the current active log information is moved offline to an archived log file, and the active log is reset. If information cannot be written to the active database log, the DBMS will halt database modifications until the log has been archived and the active log is reset.

The DBA typically controls the frequency of the log archival process by using a DBMS configuration parameter. Most DBMSs also provide a command to allow the DBA to manually request a log archival process. And remember, each DBMS performs log archival and backup differently. Refer to the sidebar “[Backing Up SQL Server Transaction Logs](#)” for a short discussion of transaction log backup considerations.

The DBA typically controls the frequency of the log archival process.

Backing Up SQL Server Transaction Logs

To ensure an efficient and effective backup and recovery strategy for your

Microsoft SQL Server databases, you will need to implement periodic transaction log backups. A transaction log backup is created using the `BACKUP LOG` command. A database can be restored to any point in time contained within the sequence of transaction logs you have backed up, up to the point of failure. If you do not back up your transaction logs before truncating them, you will be able to restore your database to only the last database backup you have created.

When Microsoft SQL Server finishes backing up the transaction log, it truncates the log's inactive portion. This frees up space on the transaction log. SQL Server can reuse this truncated space instead of causing the transaction log to continuously grow and consume more space. The active portion of the transaction log contains transactions that were still running and had not yet completed at the time of the backup.

Microsoft SQL Server will try to take a checkpoint whenever the transaction log becomes 70 percent full, or when a "log full" error occurs. Additionally, when SQL Server is shut down (unless the `NOWAIT` option is specified), it will take a checkpoint for each database.

The transaction log should not be backed up if the `TRUNC LOG ON CHKPT` database option is set to `TRUE`. If you specify `TRUNC LOG ON CHKPT` as true, Microsoft SQL Server will clear out inactive transaction log entries at every checkpoint.² This option tells SQL Server that you will not be using the transaction log for restore operations. The transaction log must still be created, though, because it is required to roll back transactions and for SQL Server to determine how to recover databases when it is restarted. Enable checkpoint log truncation only for systems where it is OK for you to lose transactions during the course of a day, because you will be able to restore your database back to only the last database backup that was taken. Applications of this nature are uncommon in most production environments.

If a transaction log has been truncated (except by a `BACKUP LOG` command), you should not back up that log until you take a database backup or differential database backup. A differential database backup will back up only the data that has changed since the last full database backup.

You should also avoid backing up transaction logs in which any nonlogged operations have occurred since the last database backup was created. If this is the situation, create a database backup or a differential database backup instead.

Finally, if any files are added to or deleted from the database, a transaction log backup should not be taken. Instead, you should create a full database backup, after which transaction log backups can be created again.

Determining Your Backup Schedule

Establishing a reasonable backup schedule requires you to balance two competing demands: the need to take image copy backups frequently enough to assure reasonable recovery time, and the need to not interrupt daily business. The DBA must be capable of balancing these two objectives based on usage criteria and the capabilities of the DBMS.

Not all data is created equal. Some of your databases and tables contain data that is necessary for the core of your business. Other database objects contain data that is less critical or easily derived from other sources. Before you can set up a viable backup strategy and schedule, you will need to analyze your databases and data to determine their nature and value to the business. To do so, answer the following questions for each database object:

Not all data is created equal.

- How much daily activity occurs against the data?
- How often does the data change?
- How critical is the data to the business?
- Can the data be recreated easily?
- What kind of access do the users need? Is 24/7 access required?
- What is the cost of *not* having the data available during a recovery?
What is the dollar value associated with each minute of downtime?

Grade each database object in terms of its criticality and volatility.

It can be helpful to grade each database object in terms of its criticality and

volatility. This can be accomplished using the grid shown in [Figure 16.2](#). The vertical axis represents a criticality continuum that ranges from easily replaceable data to data that cannot be easily replaced. The horizontal axis represents a volatility continuum that ranges from static data that changes infrequently to volatile data that changes frequently. Use this grid to diagram each database object by estimating its relative volatility and importance to the organization. Remember, these terms are somewhat vague; you will need to analyze your data and define it along the axes based on your knowledge of the data and your organization.

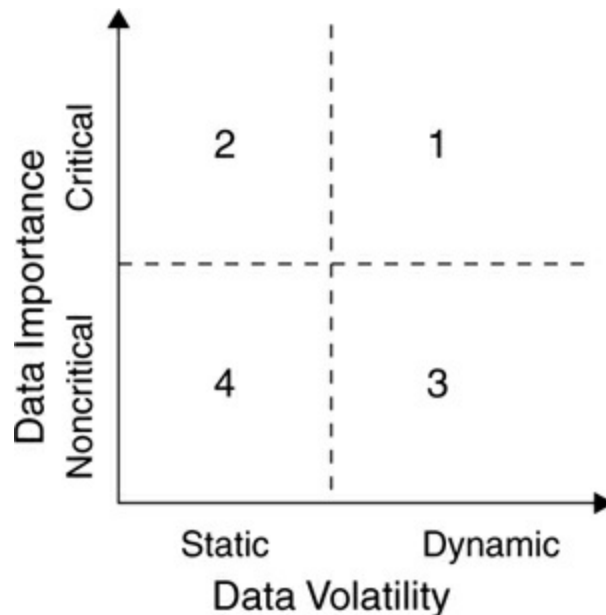


Figure 16.2. Nature and types of data

Once you have charted your database objects, you can use the diagram as a general indicator of how frequently each database object should be backed up. The DBA in charge of each application must develop the backup thresholds for each different type of data, as suggested by the grid. In general, critical data should be backed up more frequently than noncritical data, and volatile data should be backed up more frequently than static data. The key, however, is how you define the term *frequently*. For example, 1,000 updates per day might be frequent at some shops, whereas 50,000 updates per day might be infrequent at other shops. The DBA uses the grid to determine an appropriate backup schedule for each database object. The method of backup is also affected by user access needs.

Quadrant 1 on the grid identifies the critical/dynamic data in the

organization. This data is crucial to your business and it changes rapidly. As such, you must be able to recover it quickly, so you should copy it frequently. As a rule of thumb, the data should be backed up at least on a daily basis. If more than 20 percent of the data changes daily, be sure to make full rather than incremental backups.

Quadrant 2 represents critical but static data. Even though the data changes little from day to day, you will need to recover the data promptly in the event of an error because it is critical to the business. Be sure to back up this data at least weekly. Consider using incremental backups that are merged immediately upon completion to minimize the work required during a recovery.

Quadrant 3 represents volatile data that is not as vital to your business. You may be able to recreate the data if it becomes corrupted. Depending on the amount of data and the volume of change, you might not even back it up at all. For small amounts of data, a printed report may suffice as a backup. If the data fails, you could simply reenter it from the printed report. Alternatively, if data is recreated nightly in a batch job, you could simply run the batch job to refresh the data. As a DBA, you will need to ensure that the data can be recreated or copied on a regular basis. In general, more than a weekly backup for quadrant-3 data is likely to be overkill.

Quadrant 4 represents static, noncritical data. Such data does not change much and can be replaced easily. It is the least important data and should be addressed only when data in the other three quadrants has been adequately backed up. In fact, quadrant-4 data may never need to be backed up—the DBA could take a similar approach to that described for quadrant 3.

DBMS Instance Backup

In addition to being prepared for failure of individual database objects, the DBA must be prepared to recover from failure of the entire DBMS instance or subsystem. Be sure to back up all of the crucial components of the database instance, including DBMS files, system catalog and directory objects, database (archive) logs, configuration and setup files, system libraries, tape management libraries, program source libraries, and executable libraries. Of course, each DBMS and platform will have different key components that must be dealt with when planning a recovery strategy for the DBMS instance.

Be prepared to recover from failure of the entire DBMS instance or subsystem.

Recovering an entire DBMS instance is relatively rare, but it is not something that can be ignored. The DBMS instance can fail for reasons as diverse as a device failure (when critical DBMS components reside on that device), a botched version upgrade, or human error (deleting a critical file).

In certain dire situations, you may need to resort to the original installation media in order to recover a DBMS component. Of course, when restoring from the original media, you may lose any subsequent bug fixes that were applied. Therefore, the recovery process would include reapplying the maintenance as supplied by the DBMS vendor.

Recovering from a DBMS instance failure can be a very complex process. Always refer to the documentation provided with the DBMS when planning your backup and recovery strategy to make sure you are backing up everything that is necessary. Furthermore, always reread the manuals before attempting DBMS instance recovery, and never attempt to recover the DBMS instance without thoroughly understanding the reason for recovering.

Designing the DBMS Environment for Recovery

You should always take advantage of DBMS features to make recovery an easier process. For instance, most DBMSs provide the option of allocating redundant database logs. Take advantage of this option. Always allocate (at least) two copies of the active database log files to separate disk controllers, channels, and volumes. By separating the dual copies in this fashion, you are building in protection against volume failures. If one copy of the active database log is lost or damaged, operations can continue using the other copy.

Allocate redundant database logs.

Alternate Approaches to Database Backup

The backup methods discussed so far have involved copying the physical data files for a database object verbatim (or almost verbatim) to a backup device. Such image copy backups provide a fast method of backing up data and are relatively easy to use during a recovery.

The preferred method of backing up data stored in a database is to use the utilities and methods provided by the DBMS. However, other approaches can

be used, too. These approaches should be considered as special procedures to be used only under certain circumstances.

Using Database Exports to Create Logical Backups

An alternate approach to database recovery is to create an *export*, or *unload*, of the data stored in the database object. Sometimes the process of backing up just the data, and not the entire physical file, is referred to as a *logical backup*. In instances such as the following, it is quite useful to deploy logical backups:

- *Object or row recovery.* If someone inadvertently drops a table or deletes a few rows from a table, restoring them from a physical backup may be difficult. With a logical backup, the missing data is simply reloaded to the table.
- *DBMS release upgrade.* Sometimes the DBMS vendor changes the underlying database structures for a release migration. When this occurs, it can be useful to have logical backups that can be imported into the new version, instead of trying to convert existing structures—data and all—to the new format.
- *Heterogeneous database migration.* Physical data structures differ among different platforms, even when the same DBMS is being used. For example, an Oracle database on OS/390 will use different physical file structures from the same Oracle database implemented on Windows NT. Logical backups can be used to facilitate the movement of data between different databases on different platforms.
- *Data movement.* Data, once created in the database, will be shifted and moved all over the organization—perhaps to different DBMSs, perhaps to flat files, perhaps to spreadsheets. A logical backup, because it is just the data, can make it easier to move that data wherever it needs to be, whenever it must be there.

A logical backup is performed with the database up and running, so the only impact on performance will be the possible concurrent access to data by transactions and other production programs. However, as a DBA you must keep the data integrity of the logical backup in mind. Although the DBMS will use its locking mechanism to assure consistent data, referential integrity will not be guaranteed unless efforts are made to quiesce concurrent activity during the data export process.

Use logical backups to complement your physical backup strategy.

The regular creation of logical backups can complement your physical backup strategy.

Using Storage Management Software to Make Backup Copies

Storage management software can be used to back up files on managed storage devices. Some DBMSs interact with certain storage management software—for example, IBM DFSMS backups can be integrated with DB2 image copy backups. Also, EMC and other SAN vendors enable SAN duplication, which has been used by some shops to supplement database backups.

When using storage management software to back up database objects outside the scope of DBMS control, be sure to disable database write operations for all database objects that are being backed up. To accomplish this, use the DBMS commands to stop the database objects, or start them in read-only mode. Once database modification operations have been suspended, the storage management software can be used to back up the database objects. When the backup is complete, use the DBMS commands to restart the database objects in read/write mode.

Disable database write operations for all database objects that are being backed up.

Be sure you fully understand both the functionality of the storage management software and the DBMS. For example, some storage management software will not copy open files. However, files containing database data will be open if the DBMS is using them. Failing to understand such details can result in missing backups and unrecoverable data.

Additionally, when recovery is required, the storage management software should be used instead of DBMS utilities or commands. The backups were made by the storage management software and will likely be in a format that is usable only by the same software.

Document Your Backup Strategy

Once your backup strategy has been established and implemented, the backup system can run for a long time without any DBA intervention required. Such

automation is a mixed blessing, though. Over time, things can be forgotten and the DBA staff can change, both of which can cause confusion during a hectic database recovery. For this reason it is imperative that the backup and recovery strategy, implementation, and procedures be thoroughly tested and documented by the DBA.

The most important aspect of any backup plan is to thoroughly test each different type of recovery in a test environment. Be sure that you can recover from a media failure, an instance failure, and several types of application failures. Document the type of backup taken for each database object, along with a schedule of when each is backed up. Be sure that all of your databases can be recovered and that all DBAs on-site have firsthand experience at database recovery. The DBA group should schedule periodic evaluations of the backup and recovery plans for every production database.

Schedule periodic evaluations of the backup and recovery plans for every production database.

Database Object Definition Backups

In addition to regularly backing up data, the DBA should consider regularly backing up the database object definitions. Database object definitions can change over time as parameters are altered and modified. What happens if an object is inadvertently dropped and no record of the DDL used to create and change the object over time is maintained? The answer is, it can be very difficult to recreate the object exactly as it was before it was dropped.

DB2 for Linux, UNIX, and Windows provides the DB2LOOK utility to do this; Oracle offers the Export utility. If the DBMS you use does not offer such a utility, database object information can be obtained from the DBMS's system catalog. Consider generating regular reports from the system catalog that detail the specific parameters and options in effect for each database object in the production environment.

Database object definitions can change over time.

You may be able to create SQL queries against the system catalog that generate actual DDL statements that can be rerun as needed. For example, consider the following pseudo-SQL query:

[Click here to view code image](#)

```
SELECT "CREATE DATABASE ", dbname, " BUFFERPOOL ",  
bpname,  
" STOGROUP ", sname, " ;"  
FROM sysdatabases;
```

This SQL query creates SQL CREATE statements from the metadata in the system catalog. Why is this information important? Consider this tale of woe: A DBA using mainframe DB2 had to drop a partitioned tablespace and recreate it with fewer partitions. There was no way to alter the tablespace to remove partitions. After the change, there was an increased rate of locking problems. Fortunately, the DBA had saved the definition of the tablespace prior to the change. In this case, the original lock size was row, but the current lock size was page. Setting the tablespace back to row locking solved the problem—however, without that original definition, the problem would have been much more difficult to trace and resolve.

Recovery

When problems impact the database, the DBA can use the image copy backups and the database log to recover the database. Whatever the cause of the problem, the DBA must be able to recover data quickly so that the business can continue to operate. When data is unavailable, your company may be losing thousands or even millions of dollars. Recognizing the need for a database recovery is quite different from actually performing a recovery in a speedy and proper fashion. Database recovery can be a very complex task that is prone to errors and difficult to manage.

Database recovery can be a very complex task.

Recovery involves much more than simply restoring an image of the data as it appeared at some earlier point in time. A database recovery will involve bringing the data back to its state at (or before) the time of the problem. Often a recovery involves restoring databases and then reapplying the correct changes that occurred to that database, in the correct sequence.

Simply stated, a successful recovery is one where you get the application data to the state you want it—whether that state is how it was last week, yesterday, or just a moment ago. If you planned your backup strategy appropriately, you should be able to recover from just about any type of failure you encounter.

Determining Recovery Options

When a failure occurs, as a DBA you will need to ascertain whether recovery is required. If recovery is required, you will need to determine what resources (backup copies) are available and how best to perform that recovery. You need to answer several questions to determine the type and extent of failure. Your answers dictate the steps you take to recover the system.

- What type of failure has occurred: media, transaction, or database instance?
- What is the cause of the failure?
- How did the database go down: abort, crash, normal shutdown?
- Did any operating system errors occur?
- Was the server rebooted?
- Are there any errors in the operating system log?
- Are there any errors in the alert log?
- Was a dump produced?
- Were any trace files generated?
- How critical is the lost data?
- Have you attempted any kind of recovery so far? If so, what steps have already been performed?
- What types of backups exist: full, incremental, both?
- What needs to be recovered: the full database, a tablespace, a single table, an index, or combinations thereof?
- Does your backup strategy support the type of recovery required (recover to current versus point in time)?
- If you have cold backups, how was the database shut down when the cold backups were taken?
- Are all of the archived database logs available for recovery?
- Do you have a recent logical backup (EXPORT or UNLOAD)?
- What concurrent activities were running when the system crashed?
- Can you bring the DBMS instance up?
- Can you access the database objects?
- What are your system availability requirements?

- How much data must be recovered?
- Are you using raw files?

Additionally, DBMS version migration can impact recoverability. For example, consider the following sequence of events:

- A backup copy of TableA is taken while the DBMS is running on Version 10.
- The DBMS is migrated to a new version, say, Version 11.
- A problem is encountered, and TableA needs to be recovered.

Depending on the DBMS and the particulars of the new version, TableA may not be recoverable. Sometimes DBMS vendors change the format of image copy backup files, rendering any backups using the old format unusable. The same could be true for the log file—the format may have changed for a new version, rendering the old log files unreadable and the data therefore unrecoverable. Perhaps the new version changed the functionality of the RECOVER utility, such that it could not read older backup files. In either case, the DBA is stuck because there is no valid backup that can be used for recovery. For these reasons, be sure to investigate the particulars of backup and recovery for each new DBMS version before migration, and take new image copy backups after migration if the old backups are unusable.

Investigate the particulars of backup and recovery for each new DBMS version before migration.

Of course, these are just some of the questions that the DBA must be prepared to answer in order to effectively recover database objects. Additionally, the DBA will need to understand all of the details specific to the DBMS in use—for example, information on the Oracle rollback segments or how to identify a specific log RBA (relative byte address) for DB2 recovery.

When using Microsoft SQL Server, the DBA can select from several recovery models to use for database implementations. Refer to the sidebar “[SQL Server Recovery Models](#)” for more details.

SQL Server Recovery Models

SQL Server DBAs can select recovery models when creating the backup and recovery plan for their databases. A recovery model selects options for

recovery based on requirements for performance, transaction log storage, and fault tolerance. There are three recovery models: simple recovery, full recovery, and bulk-logged recovery.

Simple recovery requires less log space but has the greatest potential for data loss. With simple recovery, only full and differential backups are available for recovery.

Full recovery logs everything and trades the extra log space required for a more comprehensive approach to recovery. With full recovery, the DBA can perform point-in-time recovery using the transaction log.

Finally, *bulk-logged recovery* is optimized for bulk operations and offers the best performance of the three options. Bulk operations include BCP executions, SELECT INTO, index creation, and text updating and writing. Bulk-logged recovery does not support point-in-time recovery.

SQL Server DBAs can choose from among these models to fine-tune the backup and recovery operations of their production databases.

General Steps for Database Object Recovery

At a high level, the following steps are common to most database recoveries:

- 1. *Identify the failure.*** The detection of an outage is usually simple: Either the database is not responding to the application, or the DBMS has displayed some type of error message. Some problems are more insidious, though, such as a corrupt control file. This type of problem takes more skill to identify.
- 2. *Analyze the situation.*** The DBA must analyze the error to determine the cause, type, and scope of the failure. Based on the results of this analysis, the DBA will choose a recovery method. This is usually the most time-consuming recovery task.
- 3. *Determine what needs to be recovered.*** The DBA must determine which database objects (and perhaps other components such as logs) are failing and prepare a recovery script that is appropriate for each component. This task can also consume a significant amount of time, especially for larger systems.
- 4. *Identify dependencies among the database objects to be recovered.*** The failure of one database object can impact other database objects (e.g.,

indexes and referentially related tables). Loss of data or recovery to a prior point in time will most likely affect related database objects.

5. *Locate the required image copy backup(s)*. The closer the image copy backup is to the recovery point in time, the shorter the amount of time it will take to recover. Keep in mind other factors such as the time it takes to find tapes in the library and the possibility of the tape being at an off-site location.
6. *Restore the image copy backup(s)*. Restoration is accomplished using the database recovery utility or file system recovery command of choice.
7. *Roll forward through the database log(s)*. To recover to current or to a point in time after the image copy backup was taken, the database logs will need to be processed.

At a very basic level, every database recovery will involve most of these seven steps. However, depending on the situation and the type of recovery requested, certain steps might be eliminated or changed significantly.

Types of Recovery

Of the different types of recovery performed, the first one that usually comes to mind is a *recover to current*, to handle some sort of disaster. This disaster could be anything from a media failure to a natural disaster destroying a data center. Applications are completely unavailable until the recovery is complete. This process is depicted in [Figure 16.3](#).

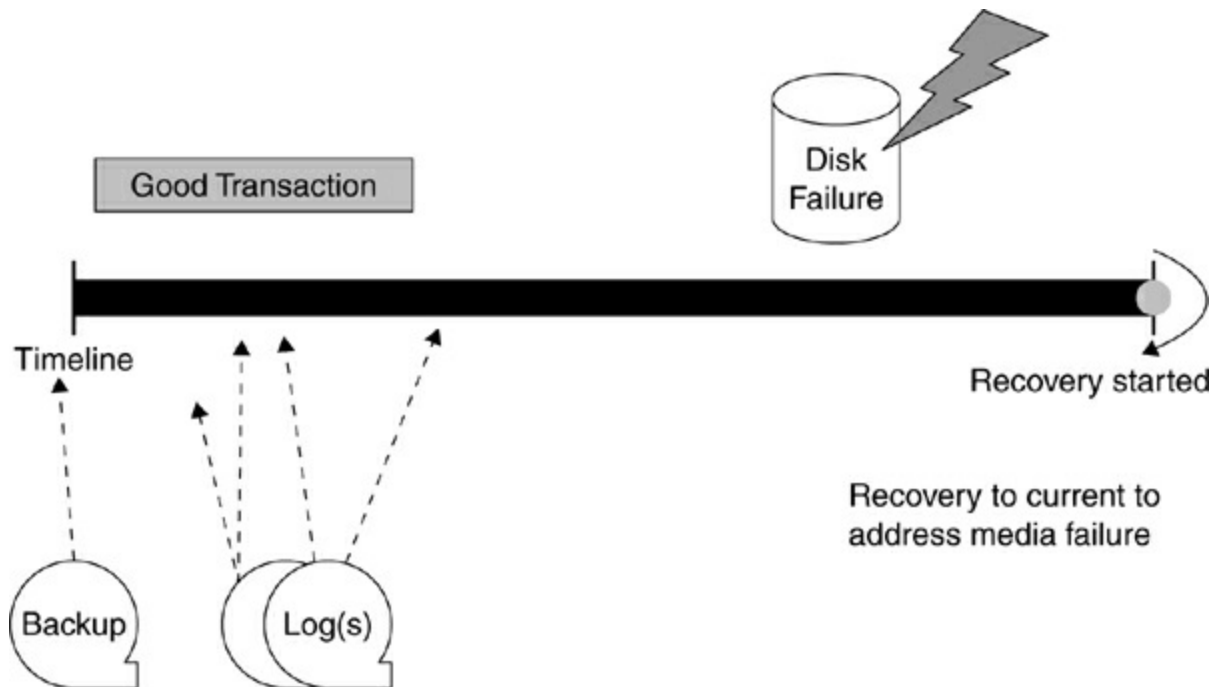


Figure 16.3. Recover to current

The first type of recovery is recover to current.

To successfully recover to current, the recovery process must be able to reset the contents of the database to the way it looked just at (or right before) the point of failure. To recover to current, the recovery process must find a valid, full image copy backup and restore that image copy. Then the recovery will roll forward through the database log, applying all of the database changes.

If the last full image copy is lost or destroyed, it may still be possible to recover if a previous image copy exists. The recovery process could start with the older backup copy, apply any incremental copies, and then roll forward through the archived and active logs. Of course, more database logs will be required in such a case, so the recovery process will take longer.

If no image copy is available as a starting point, it may be possible to recover the database object using just the database log. If the data was loaded and the load process was logged, recovery may be able to proceed simply by applying log records.

Another traditional type of recovery is *point-in-time* (PIT) recovery, which is usually done to deal with an application-level problem. PIT recovery is

sometimes referred to as *partial recovery* because only part of the existing data will remain after recovery. Recovery to a point in time removes the effects of all transactions that have occurred since that specified point in time.

Recovery to a point in time removes the effects of all transactions that have occurred since that specified point in time.

To perform a PIT recovery, an image copy backup is restored and then changes are applied by rolling forward through the database log (or log backups). However, only the log records up to the specified time are processed. Sometimes the recovery point is specified as an actual date and time; sometimes it is specified using a relative byte address on the database log. The point-in-time recovery process is depicted in [Figure 16.4](#).

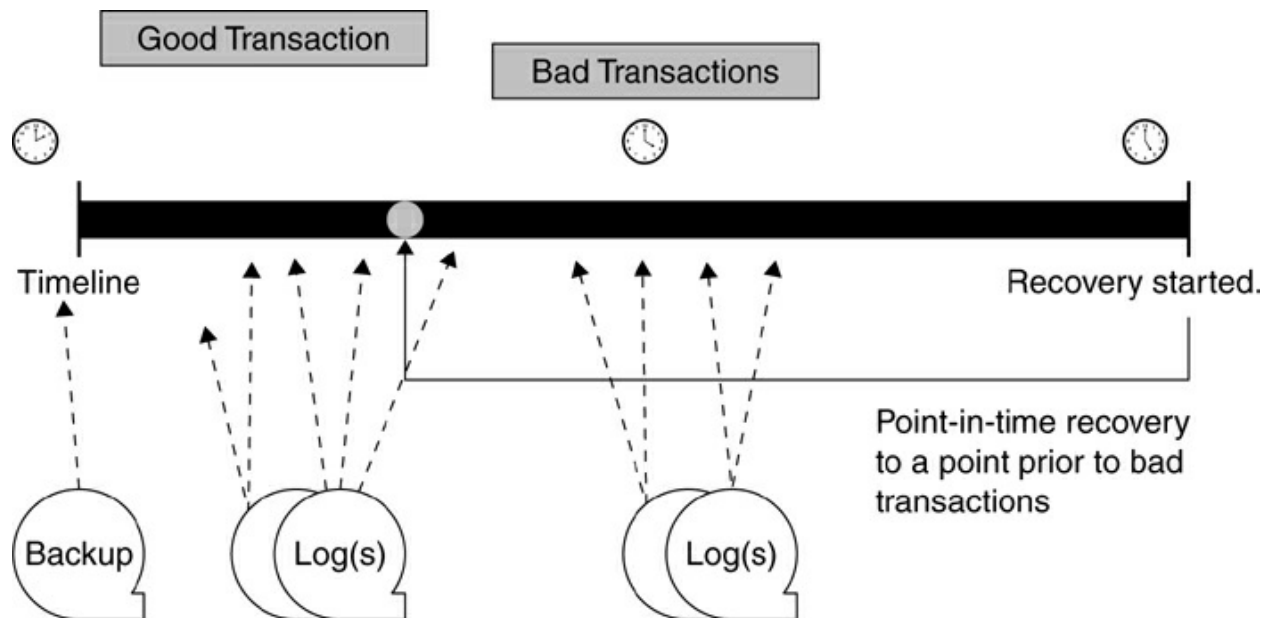


Figure 16.4. Point-in-time recovery process

To successfully recover to a point in time, the recovery must be able to reset the contents of the database to the way it looked at a prior consistent point. The key is to retain all of the good changes to the data while removing all of the “bad.” The PIT recovery depicted could be done in one of two ways, depending on the features of the DBMS and the amount of data to be recovered. It could

- Restore the image copy by rolling forward through the logs and applying the database changes up to the recovery point, or

- Not restore the image copy, instead rolling backward through the logs and removing the database changes that occurred after the recovery point

If the DBMS supports both types of recovery, the DBA should choose to deploy the one that creates less downtime. If a significant number of changes need to be removed, restoring and rolling forward usually results in less downtime. If the number of changes that must be removed are minimal, rolling backward through the logs should result in less downtime. Few DBMSs support backward log rolling, but third-party products are available that can make it a viable technique.

Regardless of the method used to perform PIT recovery, the DBA must select a *recovery point* that represents a point where the data is consistent. A consistent recovery point will assure data integrity, referential integrity, and transaction integrity. Some DBMSs provide a QUIESCE utility to achieve a point of consistency, whereas others require the database objects to be stopped (or in read-only mode) during the image copy.

An additional issue is determining what exactly has run since the bad transaction or failure from which you are recovering. The DBA can examine the job schedule, format a report from the database logs (if the DBMS provides a utility to produce such a report), and review the computer console for messages to help determine what processes ran after the failure occurred.

Transaction recovery is a third type of recovery; it addresses the shortcomings of the traditional types of recovery: downtime and loss of good data. Thus, transaction recovery is an application recovery whereby the effects of specific transactions during a specified time frame are removed from the database. Third-party software is required to perform a transaction recovery.

Third-party software is required to perform a transaction recovery.

Traditional types of recovery, both recovery to current and PIT, recover at the database object level. In direct contrast to this level of granularity, transaction recovery allows a user to recover a specific portion of the database based on user-defined criteria. This can be at a transaction or application program level. In this context, a transaction is defined by the user's view of the process. This might be the set of panels that constitute a

new-hire operation, or the set of jobs that post to the general ledger. The important point is that there may or may not be a correlation between the transactions you are trying to fix and transactions (or units of recovery) in the DBMS.

Examples of user-level transaction definitions might be

- All database updates performed by user ID DSGRNTLD since last Wednesday at 11:50 a.m.
- All database deletes performed by the application program named PAYROLL since 8:00 p.m. yesterday

Any number of problems such as the following can occur at the application level:

- Edit checks in the programs or database are not defined properly or contain bugs.
- Somebody changed the job scheduling software or didn't check for valid completion codes, so certain processes were run out of sequence, causing data problems.
- Inadequately tested code hits production.
- There are bugs in the system software.

Or maybe there is no problem, but you need to run a program multiple times against the same data. There are times, in particular in test and system test environments, when you may want to run a test and then just roll the results back and try it again (and again and again . . .).

Once you have identified the transaction to recover, you have three recovery options:

- *PIT recovery*. You can try to identify all of the database objects impacted by the application and perform traditional point-in-time recovery to remove the effects of the transactions. You would then manually rerun or reenter work that was valid.
- *UNDO recovery*. Remove only the effects of the bad transactions.
- *REDO recovery*. Remove all the transactions after a given point in time, and then redo the good transactions only.

PIT recovery was discussed earlier, so we will not cover it again here. The other two forms of transaction recovery are SQL-based application recovery techniques. The basic idea is to read the log and generate the proper SQL

statements that can be run to achieve the desired recovery results.

UNDO recovery is the simplest version of SQL-based transaction recovery

Let's first examine an *UNDO recovery*. UNDO recovery is the simplest version of SQL-based transaction recovery because it involves only SQL. To accomplish an UNDO recovery, the database logs must be scanned for the identified transaction and anti-SQL is produced. Anti-SQL reverses the affect of SQL by

- Converting inserts into deletes
- Converting deletes into inserts
- Reversing the values of updates (e.g., UPDATE "A" to "X" becomes UPDATE "X" to "A")

Once the anti-SQL is generated, it is run using an interactive SQL script to perform an UNDO recovery, which is shown in [Figure 16.5](#). The portion of the database that does not need to be recovered remains undisturbed. When undoing erroneous transactions, recovery can be performed online without suffering an outage of the application or the database. Indeed, the DBMS must be up and running in order to execute the anti-SQL. However, certain applications may need to be brought down for the duration of the UNDO recovery to eliminate the potential for data anomalies causing additional failures.

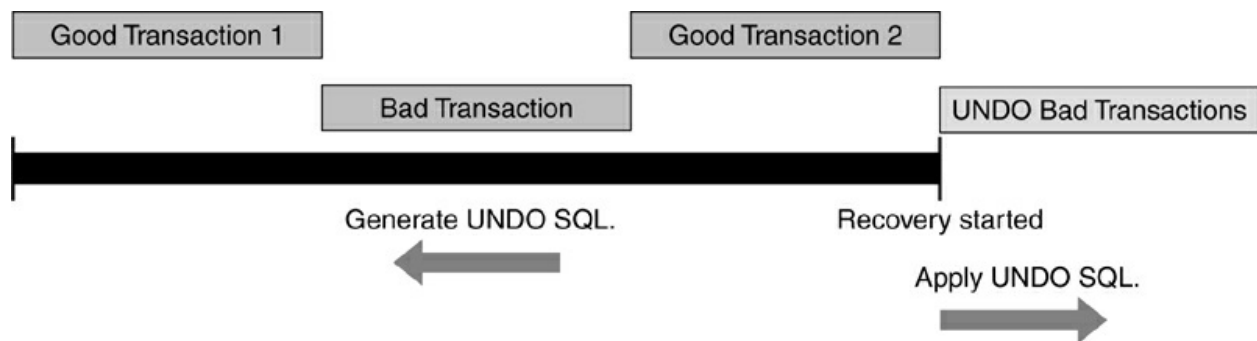


Figure 16.5. UNDO recovery process

A REDO recovery is a combination of PIT recovery and UNDO recovery, with a twist.

A REDO recovery is a combination of PIT recovery and UNDO recovery,

with a twist. Instead of generating SQL for the bad transaction that we want to eliminate, we generate the SQL for the transactions we want to save. Then we do a standard point-in-time recovery to eliminate all the transactions since the recovery point. Finally, we reapply the good transactions captured in the first step. This process is shown in [Figure 16.6](#).

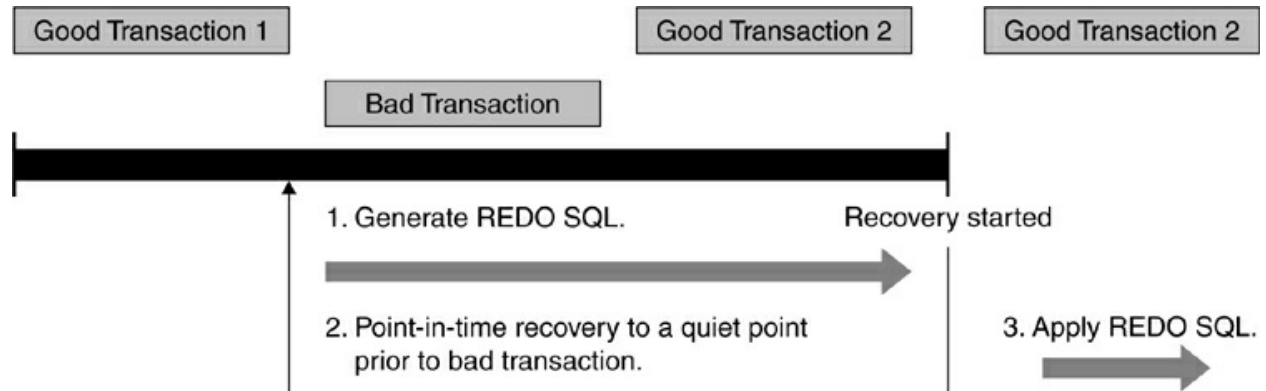


Figure 16.6. REDO recovery process

Unlike the UNDO process, which creates SQL statements designed to back out all of the problem transactions, the REDO process creates SQL statements that are designed to reapply only the valid transactions from a consistent point in time. Since the REDO process does not generate SQL for the problem transactions, performing a recovery and then executing the REDO SQL can restore the database object to a current state that does not include the problem transactions.

When redoing transactions in an environment where availability is crucial:

1. Perform a recovery to a point in time.
2. Bring the application and database online.
3. Redo the subsequent valid transactions to complete the recovery. This step must be done while the database is online for concurrent read/write operations.

If your approach to backup and recovery is to use *storage management software*, individual image copy backups for database objects will not exist. In this case, you will need to deploy storage management software to perform the recovery. The actual recovery process will depend on the type of storage management software in use, as well as how (or if) it interacts with the DBMS recovery mechanisms.

Off-site disaster recovery is the rarest, but most comprehensive, type of

database recovery. An off-site disaster recovery is required when a natural disaster (fire, flood, etc.) or other severe accident makes it impossible to use your primary data-processing center. In this case, you need to reestablish your entire system environment, as well as recover your DBMS, database objects, and data. Disaster recovery is discussed in more detail in [Chapter 17](#), “[Disaster Planning](#).”

Off-site disaster recovery is the most comprehensive type of database recovery.

Choosing the Optimum Recovery Strategy

So, what is the best recovery strategy? It depends.

Historically, recovery was performed mostly for disasters and hardware failures, but this is simply not the case anymore. In fact, the majority of recoveries these days result from application problems. Recent industry analyst studies have shown that most system downtime is caused by software problems—not hardware problems.

In reality, very few DBAs ever need to perform a true disaster recovery except during tests. Though media continues to fail, it fails relatively infrequently these days. User errors and application failures are the *most common* reasons for database recovery and therefore the primary causes of system unavailability.

User errors and application failures are the most common reasons for database recovery.

Software problems and bugs might cause only certain transactions to be in error and in need of repair. Unfortunately, as databases grow in size and complexity, so do the chances that bad transactions will corrupt the data on which your business depends. Transaction recovery may seem like the answer to availability problems, but there are a number of cases where transaction recovery is neither possible nor advisable. In determining the type of recovery to perform, the DBA must consider several questions:

- *Transaction identification.* Can all the problem transactions be identified? You must be able to actually identify the transactions that will be removed from the database for transaction recovery to work. Can all the work that was originally done be located and redone?

- *Data integrity.* Has anyone else updated the rows since the problem occurred? If so, can you still proceed? Is all the data that is required still available? Intervening reorganizations, loads, or mass deletes can require the use of an image copy backup, thereby eliminating UNDO recovery. Will the recovery cause any other data to be lost? If so, can the lost data be identified in some fashion and reapplied?
- *Speed.* If multiple techniques are viable, which one is likely to perform the fastest? How many database logs are required to perform the recovery? Can anything be done to reduce the number of logs, such as merging incremental copies?
- *Availability.* How soon can the application become available again? Can you afford to go offline?
- *Invasiveness.* How invasive was the failure to your database? Were decisions made based on bad data? Can *any* subsequent work be trusted?

What is the cost of rework?

All of these questions actually boil down to a question of cost. What is the cost of rework, and is it actually possible to determine what would need to be redone? This cost needs to be balanced against the cost of long scans of log data sets to isolate data to redo or undo and the cost of applying that data via SQL. Of course, an additional question looms large: Which of these recovery techniques are actually available at your site and will they work for the DBMS in question?

Many factors influence the duration of the recovery process. The DBA can implement measures to reduce downtime by developing a smart backup and recovery plan. The following factors can shorten the duration of a recovery:

- The smaller the size of the components that need to be recovered, the shorter the recovery process will be. In general, the less you have to do, the less time it will take.
- Consider partitioning database objects and backing up and recovering at the partition level. Sometimes a failure that would otherwise impact an entire database object can be limited to impacting only a single partition.
- Consider keeping image copy backups and log archive files on disk.

Because disk file access is quicker and processes do not need to wait for tape mounts, using disk instead of tape can speed up the recovery process.

- Test your image copy backups to make sure they are valid. Encountering an invalid image copy during the recovery process will lengthen the duration of recovery. When invalid image copy backups are found, steps can be taken to create a new, valid image copy backup before causing a negative impact on recovery.
- Automate your backup and recovery procedures to the greatest extent possible. Automated procedures remove manual error from the equation, thereby minimizing downtime.
- Whenever possible, design databases with as few dependencies as possible. Autonomous database objects can minimize the duration of a recovery because fewer related database objects need to be recovered at the same time.
- Finally, be sure that all DBAs understand the recovery procedures for each database object under their control.

Matching Type of Failure to Type of Recovery

Matching the type of failure to the appropriate type of recovery is good practice. Of course, there are exceptions to every rule, so the following generalizations should be treated as a starting point only.

Match the type of failure to the appropriate type of recovery.

Recovering from a *media failure* usually involves a recover to current. When the media fails, the database objects residing on that media most likely will not be able to be accessed or changed. The general desire in such a situation is to recover all database objects on the failing media to the point just before the failure—in other words, the DBA will try to recover all activity and data for these database objects.

Recovering from a *transaction failure* usually involves a point-in-time recovery or a transaction recovery. By definition, a transaction recovery is caused by an erroneous or incorrect execution of a program. The database changes resulting from the improperly run program must be removed from all database objects affected.

Recovering from a *database instance or subsystem failure* will most likely involve a recover to current. The goal of such a recovery is to bring the data in all database objects within that instance or subsystem back to the way it was before the point of failure, and in a consistent state.

Index Recovery

Recall from the previous discussion on image copy backups that there are two options for index recovery:

- Rebuilding the index from the table data
- Recovering the index from a backup copy of the index itself

Most DBMSs provide the capability to rebuild an index from the table data. Some DBMSs support both methods of index recovery. When both methods are available, the DBMS may require you to choose one or the other method for each index that you create. As a DBA, you will need to examine the index recovery methods available to your DBMS and choose the appropriate method for each index.

In general, the larger the amount of data to be indexed, the bigger the actual index will be and the longer it will take to rebuild that index from the data. Therefore, consider using image copy backup and recovery for indexes when the database object is very large. Of course, the definition of “very large” is vague and will differ from organization to organization depending on your database usage and the amount of downtime your company can tolerate.

When you choose to use backup and recovery for indexes, be sure to synchronize the backup and recovery with the indexed database objects. In other words, when you make a backup of a tablespace, be sure to also copy any associated indexes. Failure to do so can result in data integrity problems.

When you make a backup of a tablespace, be sure to also copy any associated indexes.

Testing Your Recovery Plan

Every DBA should prepare a recovery plan for each database object and test it frequently. A recovery plan describes the procedures you will use to recover in the event of hardware failure or a local site disaster, such as fire. You should develop a recovery plan and test it often (ideally no less than

twice per year).

Develop a recovery plan and test it often.

To develop your recovery plan:

- Write out all aspects of the recovery plan in detail, documenting each step.
- Include all the scripts required to back up and recover each database object.
- Review the plan with everyone who may be called on to implement it.
- Include a contact list with names and phone numbers of everyone who may be involved in the recovery.
- Keep the recovery plan up-to-date by modifying it to include every new database object that is created.

Testing your recovery procedures on a regular basis includes running regular recovery tests for individual database objects, databases, and the entire database system. By testing recovery procedures, the DBA assures that the backup and recovery scripts work and that every database object is indeed recoverable. Furthermore, regular recovery testing is on-the-job training for the inevitable, stress-filled production database recovery. DBAs who test their recovery plan will be more familiar with the tools and scripts needed to perform recovery.

Recovering a Dropped Database Object

Until now, we have focused on backing up and recovering data. However, it is also possible to drop a database object unintentionally. When such an error is recognized, the DBA will need to recover the database object as quickly as possible to avoid availability and integrity problems.

Recovering a dropped object requires extra steps beyond a normal recovery. Depending on the DBMS and the tools available, it can sometimes be very complicated.

Recovering a dropped object requires extra steps beyond a normal recovery.

Typically, each DBMS identifies the database objects under its control by an internal identifier. When an object is dropped and recreated, the internal

identifier for that object usually will change. Therefore, recreating the object using the same DDL and running a recovery using a prior image copy backup usually will not work.

To recover a dropped database object, the DBA may need to translate the internal identifier of the old database object to the internal identifier of the new database object. Refer to the sidebar “[DB2 and DSN1COPY](#)” for an example. To accomplish this, you must recreate the dropped database object using the exact same definition that was in effect when the image copy backup was taken. If a column definition differs between the dropped object and the newly recovered object, the recovery will probably fail.

Of course, the preceding scenario does not apply if you have a logical backup file. If that is the case, you should be able to recreate the database object and set up a load script to repopulate the data.

At any rate, regardless of the method used to recover the dropped database object, keep in mind the ramifications of dropping objects.

DB2 and DSN1COPY

DB2 for z/OS supplies a utility called DSN1COPY to enable database files to be copied outside the control of the DBMS. DB2 does not need to be operational for DSN1COPY to be able to run.

Furthermore, the DSN1COPY utility can be used to modify the internal object identifier used by DB2 to identify tablespace and index space data sets. Using the OBIDXLAT option, the DBA can direct DSN1COPY to translate one object identifier to another so that data can be recovered to similar (or restored) database objects, even if the internal identifier does not match the identifier in the image copy backup.

When a database object is dropped, all subordinate database objects are also dropped, as are security authorizations, referential constraints, and check constraints. Once you have recovered a dropped object, you need to address these other issues.

Recovering Broken Blocks and Pages

A broken block or page is a section of a tablespace or index that contains bad or inconsistent data. Data may be inconsistent due to a broken or orphaned chain, referential constraint violations, a damaged recovery log, a missing or

extra index entry, or some other arcane problem. To recover an index with a broken page you can simply rebuild the index from the data in the tablespace.

Rebuild the index from the data in the tablespace.

Tablespaces are a different proposition. Sometimes simply stopping and starting the tablespace or recycling the DBMS instance can fix a broken page. Additionally, some DBMSs come with a repair utility that can be used to pinpoint locations within a file based on offsets and replace data at the bit or byte level. Before using any such repair tool, be sure to completely read the DBMS instruction manuals. Furthermore, be sure that the corrective action will actually fix the broken page. Repair utilities can be invasive and damaging to the contents of the database. Sometimes the DBMS vendor recommends that such activities be performed only under the direction of the DBMS vendor's technical support staff. Heed these recommendations to avoid causing further damage to your data.

Once you have repaired the information, you may need to recover the tablespace to current. The recovery will apply the log records that recorded the activity that caused the damage.

Populating Test Databases

One of the best ways to populate a test environment is with a production database backup. Such a database backup contains a valid set of data from the production environment at a given point in time. To create the test environment, the DBA can use the RECOVER utility or perhaps an UNLOAD utility that can read backup files.

Populate the test environment with a production database backup.

Creating a test system from a production database can impose specific requirements on the structure and timing of the production backup, as well as the structure of the test environment. If the production database backup is not taken during a period of no activity, the restore to the test environment will require the production log files as well.

Populating the test environment with production data can be a feasible approach to resolving some types of production problems. For example, consider a 24/7 application with month-end reports and month-end database

backups. On the tenth day of the month, a business user finds a problem in the month-end report. To resolve the issue, the programmer needs to know the contents of the database as of the first of the month. However, the database has changed significantly over the past ten days. Restoring the production database with month-end data is not an option—business transactions must continue. In such a scenario, the DBA could use the month-end production database backup to populate the test environment with month-end data. The programmer can now fix the month-end report by using the test database with the month-end production data.

Alternatives to Backup and Recovery

The creation of image copy backups for recovery is the most common and reliable method of insuring against failures and lost data. However, several alternatives exist that either augment or perhaps replace standard backup and recovery methods. The next few sections briefly examine some of these alternatives.

Standby Databases

Oracle introduced the concept of a [*standby database*](#) back in Version 7. A standby database is an identical copy of an online production database that is close to being up-to-date in terms of its data content. The standby database may not be 100 percent up-to-date because of system latency between applying updates from the online production database to the standby database. When a failure occurs, control is transferred to the standby database, which is then opened as the online production database to allow normal activity to continue.

Usually, a standby database is created initially by restoring a cold backup. Then, all the archive logs from the production database must be copied to the standby database and applied. For all intents and purposes, the standby database is continuously running in recovery mode because it is applying the archive logs from the production database as operations are performed on the production database.

But, beware. A standby database does not eliminate the need for normal backups to be performed on the production database. Operations performed in error on the production database will be duplicated on the standby database. For example, a row deleted accidentally will also be deleted on the

standby machine.

A standby database is fairly easy to implement and provides quick recovery from certain types of failures. However, standby databases are almost never completely up-to-date: If a production archive is lost, the standby database becomes unsynchronized, and only an entire database can be duplicated (instead of just certain database objects).

A standby database is fairly easy to implement.

Standby databases can be ideal for disaster recovery.

Replication

Data replication involves storing and maintaining redundant data in a separate copy of the database. Replicated data can be a subset of the rows and/or columns of the original database. A simple form of replication can be implemented by copying entire tables to multiple locations at various times throughout the day. Of course, this is easier said than done. Some DBMSs provide automated replication features, and for those DBMSs without such features, third-party tools are available.

Some DBMSs provide automated replication features.

There are two basic technologies for replicating data among databases:

- Snapshot replication
- Symmetric replication

Snapshot replication produces a copy of database tables on a target system based on a query of the source database. At the time a snapshot is initially taken, the specified query (maybe an entire table) is run, and the resulting data is loaded into the target snapshot table. Furthermore, each replica database should contain accurate, up-to-date information. When multiple replicas exist, the DBA must facilitate the update of all replicated copies at the same time. If replicas are updated on different schedules, it becomes burdensome to keep track of the state of each replica.

The advantage of snapshot replication is its ease of implementation. However, snapshot replicas can become out-of-date very rapidly, and refreshing the replica can cause administrative and performance problems.

Symmetric replication is a more robust implementation of replication

because it keeps the replicas up-to-date. Symmetric replication can be set up to ensure that no transaction is fully committed until all the modifications have been committed locally for all replicas. Alternatively, it can replicate asynchronously, allowing each database node to run at full speed without holding up local updates because of remote database speed issues. The updates are applied later, after the COMMIT is taken for the master database.

The biggest advantage of symmetric replication over snapshot replication is the automatic synchronization of modifications from the master database to the replicas. However, symmetric replication is more difficult to set up and administer, high transaction volumes can cause performance degradation, and a network failure can cause the database to fail because modifications cannot be synchronized.

Replication is not a substitute for backup and recovery, but it can be helpful in some situations. With snapshot replication, it may be possible to use the snapshot replicas as a point in time before a failure. However, symmetric replication has the same problem as a standby database—the problems would be propagated to the replica. Both types of replication may be useful for a disaster recovery.

Replication is not a substitute for backup and recovery.

Disk Mirroring

Mirroring disk devices can add an extra level of protection to databases. *Disk mirroring* occurs by allocating a secondary device that will contain a duplicate copy of the primary device. All data modifications made to the primary device are also made to the mirrored device. If the primary device fails, the mirrored device can be used without causing a system outage. Of course, disk mirroring consumes double the amount of disk storage for the same amount of data.

Disk mirroring differs from replication in that the data duplication occurs at the device level instead of the database level.

Disk mirroring can remove the need to recover from a media failure.

Once again, disk mirroring does not replace the need for backup and recovery, because the mirrored data will be the same as the data that is

experiencing the problem. Disk mirroring, though, can remove the need to recover from a media failure. Instead of recovering database objects on the failing disk, the DBA can switch to the mirrored data.

Summary

An in-depth backup and recovery plan for every database object is an integral part of a database implementation plan. It is the duty of the DBA to ensure that every piece of critical data in the database is protected and can be recovered if a problem arises. Furthermore, the DBA must be able to minimize the amount of downtime associated with recovering from database integrity problems and failures. The longer the database is down, the larger the financial impact will be to the business.

Review

1. What is the difference between a logical backup and an image copy backup?
2. Why is a quiesce point important for a point-in-time recovery?
3. Why can't standby databases, replication, or disk mirroring be used to replace traditional backup and recovery?
4. What factors determine whether a full or incremental image copy backup should be taken for a database object?
5. Name and describe the three types of database failures that may require recovery.
6. Describe two ways to recover an index.
7. What is meant by the term *log archival*, and why is it important to assure recoverability?
8. Name four factors that impact the duration of a recovery.
9. The DBA should grade each database object in terms of its _____ and _____ to determine how frequently it should be backed up.
10. Name the different types of recovery and discuss the factors that influence when each type of recovery should be performed.

Suggested Reading

- de Guise, Preston. *Enterprise Systems Backup and Recovery: A Corporate Insurance Policy*. Boca Raton, FL: Auerbach/CRC Press (2009). ISBN 978-1-4200-7639-4.
- Freeman, Robert, and Matthew Hart. *Oracle RMAN 11g Backup and Recovery*. Berkeley, CA: Oracle Press/McGraw-Hill (2010). ISBN 978-0-07-162860-0.
- Gu, Lijun, et al. *DB2 UDB Backup and Recovery with ESS Copy Services*. San Jose, CA: IBM (2002). ISBN 0-7384-2514-1.
- McBath, Frank. *SQL Server Backup and Recovery: Tools and Techniques*. Upper Saddle River, NJ: Prentice Hall (2001). ISBN 978-0-13-062298-3.
- Preston, W. Curtis. *Backup & Recovery*. Sebastopol, CA: O'Reilly (2007). ISBN 978-0-596-10246-3.
- Velpuri, Rama. *Oracle Backup and Recovery Handbook*. Berkeley, CA: Osborne/McGraw-Hill (1995). ISBN 0-07-882323-4.

17. Disaster Planning

A disaster recovery plan is like insurance—you're glad you have it, but you hope you never need it. With automobile insurance, you pay a regular fee so that you are covered if you have an accident; in other words, it's an investment. A disaster recovery plan is similar in that you invest in it by designating a disaster recovery site, shipping backup copies of the data off-site, preparing recovery jobs, and practicing the recovery procedures.

Invest in a disaster recovery plan.

The Need for Planning

Disaster recovery planning, also called *contingency planning*, is the process of preparing your organization's assets and operations in case of a disaster. But what is a disaster? SunGard Recovery Services (1995) provides a good definition of *disaster*: any unplanned, extended loss of critical business applications due to lack of computer processing capabilities for more than a 48-hour period. Your own definition may be more or less stringent with regard to the time frame, but the basic definition is a sound one.

The *DB2 Developer's Guide* (Mullins 2012) defines a disaster as “any event that has a small chance of transpiring, a high level of uncertainty, and a potentially devastating outcome.” This, too, is a workable definition for a disaster.

Most of us have witnessed a disaster, at least on television. Floods, earthquakes, hurricanes, and fires are some examples of natural disasters. Disasters can be man-made, such as electrical failure, bursting pipes, and war. Many of us have had our basements flooded or been in an automobile accident. A disaster does not have to have global consequences in order for it to be a disaster for you.

You must recognize potential disasters and understand their consequences. How these disasters might impact your business is the sole purpose of disaster recovery planning. If your business is on a coast, the likelihood of hurricanes and floods increases. If your business is located in the Midwest, you have a better chance of experiencing tornadoes. If your business is located in the North, blizzards and severe cold weather will pose more of a

risk. California businesses are more apt to worry about earthquakes.

Even though disasters are unpredictable and unlikely, every organization should have a comprehensive and tested plan to cope with a disaster situation. For example, consider some recent disaster situations. In 2005, Hurricane Katrina impacted the southern Gulf region of the United States. It is the costliest natural disaster, and one of the five deadliest hurricanes, in the history of the United States. Hurricane Katrina damaged or destroyed many businesses, including oil platforms, refineries, and a large portion of the forestry industry in Mississippi. Furthermore, hundreds of thousands of local residents were left unemployed. And, of course, we cannot forget the 9/11 terrorist attacks on the World Trade Center in 2001 which had severe consequences for any business with a data center in one of the buildings. Furthermore, many disasters are not location specific. Sabotage, computer viruses, vandalism, air conditioning or heating failures, and health or environmental hazards can happen anywhere on the planet.

Just because your organization has not yet experienced a disaster, or is not in a high-risk area, does not absolve you from the need for contingency planning—especially for your databases. In the wake of a disaster, companies with a disaster plan will be able to service their customers again much more quickly than those companies without one. Indeed, a company facing a disaster without a disaster recovery plan may never resume business.

Database disaster recovery must be an integral component of your overall business recovery plan. A disaster recovery plan must be global in scope. It must handle business issues such as alternate locations for conducting business, communication methods to inform employees of new locations and procedures, and publicity measures to inform customers how to transact business with the company post disaster. It must restore the IT infrastructure. Finally, and most important to our discussion, a component of that plan must be for the recovery of database and DBMS operations.

Database disaster recovery must be an integral component of your overall business recovery plan.

However, to what extent should a company take disaster planning? Before your company can ascertain the appropriate level of recoverability, you must analyze the risks and determine the objectives.

Risk and Recovery

The goal of a disaster recovery plan is to minimize the costs resulting from losses of, or damages to, the resources or capabilities of your IT facilities. The success of any database disaster recovery plan depends a great deal on being able to determine the risks associated with data loss. What is the impact to your business if the data is lost?

Evaluate each database object for disaster recovery.

As with local database recovery, the DBA must perform an evaluation of each database object for disaster recovery. Recall from [Chapter 16 \(Figure 16.2\)](#) the grid we used to evaluate database objects for criticality and volatility. Although the emphasis in disaster recovery is more on criticality than volatility, the dynamism of the data still plays a role. Although it is quite likely that each database object will be graded similarly for disaster recovery purposes, it is still a good practice to reevaluate each database object in terms of the consequences of its complete loss in a disaster.

When analyzing the risk associated with losing data, keep in mind the different types of risk. With regard to data, there are three categories of risk—financial loss, business service interruption, and legal responsibilities—and varying degrees of risk within each category. The unavailability of each application has a different impact on the company's bottom line. Consider a brokerage firm, for example. Having the stock-trading system unavailable would cause a greater loss than having the employee payroll application unavailable, not only because current business operations could not be conducted in a timely manner, but because customers would stop trusting the brokerage with future business. And if a major stock exchange such as NYSE or NASDAQ were to incur an outage, the impact would be even greater—affecting the global economy.

Business needs must dictate your priorities.

As you create your database disaster recovery plan, remember that business needs must dictate your priorities, and not technical needs and issues. Consider separating your systems into critical and noncritical applications based on business needs. The task of defining *criticality* can become a political nightmare if it is left to each business unit. The decision to rank one system as more critical than another must be made at a high level—

with the overall business in mind—and not individually by each business unit.

Rank your applications into groups to determine which applications have the biggest impact if they are down.

It is a good idea to rank your applications into the following groups to determine which applications have the biggest impact if they are not available:

- *Very critical applications.* The most critical applications in your organization will require current data upon recovery. These applications are the most difficult to develop a disaster plan for because more steps are required for off-site backup and recovery. Additionally, the most critical applications in your shop will be the first ones that must be made operational at the recovery site. Try to limit the number of applications designated as very critical; any more than a half-dozen or so will be unmanageable in a disaster situation.
- *Business-critical applications.* Business-critical applications are important to your organization and should be the next group to recover after the very critical applications. A business-critical application frequently requires current data, but it may not be available at the remote site within the first couple of days. For example, consider the applications for a telephone service provider. The system that delivers phone service would be a very critical application; customer billing would be a business-critical application.
- *Critical applications.* A critical application differentiates itself from a business-critical application by its immediacy or data currency needs. This group of applications, though important, need not be available immediately. However, if the disaster persists for a week or longer, the business requires the application. Critical applications should not be recovered until the very critical and business-critical applications are up. The requirements of these applications vary from up-to-date to possibly day-old or week-old data.
- *Required applications.* Required applications are not critical but must be backed up such that they can be recovered at the remote site if needed. Data from the last available backup is usually sufficient to support such applications.

- *Noncritical applications.* Noncritical applications need not be supported in the event of a disaster. Very few applications fall into this category—if the application is not critical, why was it developed in the first place?

The criticality of an application must be based on the overall importance of the application to the organization. Factors that impact criticality include the financial cost of the application being unavailable, the legal responsibilities your organization has for each application and its data, and the cost of losing customers and future business due to a delay in bringing the application back online.

As a DBA, you must create the disaster recovery plans with application criticality in mind. In this way, the most critical data can be recovered and made available immediately in the event your company experiences a disaster. Based on these application rankings, an appropriate backup strategy can be deployed for each database object supporting your database applications.

General Disaster Recovery Guidelines

During a disaster recovery, your goals are to minimize downtime and loss of data. Whether you achieve these goals is primarily determined by the preparations you have made.

Minimize downtime and loss of data.

Planning for disaster recovery is an enterprise-wide task. Remember that DBMS and database recovery is just one component of an overall disaster recovery plan. When your organization creates a disaster recovery plan, it needs to look at all of its business functions and operational activities—customer interfaces, phone centers, networks, applications, and every company function that can be impacted by a disaster. However, this chapter addresses only DBMS- and database-related recovery issues. For a comprehensive discussion of disaster recovery, consult the books listed at the end of the chapter.

The Remote Site

When a disaster strikes, you will need an off-site location where you can set up your company's operations. The site must be located far enough away

from your primary site so that a natural disaster will not disrupt both sites. Sometimes the difference between success and failure is simple common sense. For example, your remote location should not be on the same power grid, in the same floodplain, or along the same earthquake fault line as your primary site. Your company may even select several sites for different corporate functions.

If your company is large enough to have more than one data center, you may be able to use each site as a backup for one of the others. For example, if your Houston data center is destroyed, you can move operations to your Pittsburgh data center.

Other companies set up a specific remote location where they send their data on a regular basis. In the event of a disaster, they simply connect to the backup system. This is an expensive alternative, but necessary for some businesses.

Yet another approach is to sign up with a disaster recovery service provider. The service provider maintains the equipment necessary for you to recover your operations on its computers in the event of a disaster. In this case, you are responsible for ensuring that the proper applications and data are available to be installed at the service provider's location if a disaster strikes your site.

Storage of backup materials is another issue. Ideally, they would be stored for safekeeping at the recovery site, but if this is not possible, another off-site storage location should be designated. If a disaster occurs, you will need to provide a mechanism to move the recovery materials from the storage location to the recovery location.

Backup materials must be stored for safekeeping.

The Written Plan

A written plan is the foundation of any good disaster recovery plan. The plan should be distributed to all key personnel in the disaster recovery scenario. Each participant should keep a copy of the plan at home as well as at the office. A copy of the disaster plan should be kept at the recovery site as well.

Perhaps the biggest challenge to implementing a successful disaster recovery plan is keeping the plan current and relevant. Maintain the plan as a part of your everyday activities. Be sure that your DBA procedures

automatically include disaster recovery plan updates. For example, whenever you create a new database object, be sure to incorporate that object into the disaster recovery plan. Likewise, whenever a database object is dropped, remove it from the disaster recovery plan. Furthermore, whenever you add a new application to the system, be sure to evaluate the criticality of the application and include it in the disaster recovery plan.

Simply stated, the disaster recovery plan is a living document that will change as your systems, requirements, and usage change. Whenever the plan changes, be sure to destroy all of the outdated copies of the plan and replace them with the new plan.

The disaster recovery plan is a living document.

Writing out the specific procedures and policies to follow for an off-site disaster recovery has several benefits:

- It causes you to formulate the explicit actions to be taken in the event of a disaster.
- It makes you order these actions into specific sequential steps.
- It forces you to be specific about the tools to be used and the exact backup information required.
- It documents the location where all the required information is stored and how it is to be made available at the recovery site.
- It provides a blueprint for others to follow, in case those who are most familiar with the plan are not available.

Another benefit of creating a detailed disaster recovery plan is that it will highlight areas that need to be improved within your backup procedures and organization.

Be sure to include all of the interested parties in the disaster recovery planning process. These include not just DBAs and systems programmers, but also end users, system operators, business managers, and perhaps even your customers. Your disaster recovery plan should include the following sections:

- *Off-site location.* List the address of the remote location(s), along with the phone number, fax number, and address of the contact at each remote site. Additional useful details could include a list of nearby

hotels, options for travel to the recovery site, details of how expenses will be handled, and other pertinent information.

- *Personnel*. List the name and contact information for each member of the recovery team. Be sure to include the work, home, and mobile phone numbers for each team member.
- *Authorizations*. List the security authorizations necessary for the recovery operations and the personnel to whom they've been granted.
- *Recovery procedures and scripts for all system software, applications, and data*. Be sure to provide the complete step-by-step procedures for the recovery of each piece of system software, every application, and every database object, and the order in which they should be restored. Part of this section should be a listing of all the installation tapes for system software as well as the tapes for all maintenance that has been applied. Options for database recovery procedures will be covered later in this chapter.
- *Reports*. List the reports you will need at the recovery site to ensure a complete recovery. The reports should list each backup tape, its contents, when it was produced, when it was sent from the primary location, and when it arrived at the remote site. As an additional component, include a description of the naming conventions for the remote site backup files.

The disaster recovery plan should include instructions for a complete off-site disaster recovery. However, before you commit your disaster recovery procedures to paper, you need to make a number of decisions. The first decision to be made is to prioritize your disaster recovery goals. Do you want to get the system up as quickly as possible? Or is it more important to lose as little data as possible? Or perhaps your most important goal is to avoid reprocessing data as much as possible. The disaster recovery plan should be written in accordance with your specific recovery goals.

The disaster recovery plan should include instructions for a complete off-site disaster recovery.

Testing Your Disaster Plans

Once the disaster recovery plan is written, be sure to schedule regular tests. It is a good practice to test your disaster recovery plan at the remote recovery

site at least once a year. You should also consider testing the plan after the following events:

Test your disaster recovery plan at the remote recovery site at least once a year.

- Significant change in daily operations
- Change in system hardware configuration
- Upgrade of the DBMS (or related system software)
- Loss (or hire) of personnel responsible for the recovery
- Move of the primary data center to a new location
- Change in daily backup procedures
- Addition of major new applications or significant upgrades of existing critical applications
- Major increase in the amount of data or the number of daily transactions

Use a disaster recovery test to discover weaknesses and errors in the plan. After the test, be sure to update the disaster recovery plan to address the problems. A valid disaster recovery test need not end in a successful recovery—although that is the desired result. A disaster recovery test that reveals weaknesses in the plan serves a useful purpose.

Another consideration for scheduling regular disaster recovery tests is to assure the readiness of your personnel. The best way to prepare for a disaster is to practice disaster recovery. The process of actually implementing the plan forces you to confront the many messy details that need to be addressed during the recovery process. Testing can help to identify issues such as those described in the sidebar “[A Lengthy Outage](#).” Testing also helps you to become familiar with the tools and procedures you will use during an actual disaster recovery.

Regular disaster recovery tests assure the readiness of your personnel.

Actually, a scheduled test of the disaster recovery plan is probably a poor idea. A disaster recovery test should work more like a pop quiz that doesn’t give you the opportunity to prepare. One day your boss should come to work

and announce that the building was just destroyed. Who should be called? Is everyone available? How can you get the right people to the remote site for recovery? Can you get your hands on the disaster recovery plan?

A Lengthy Outage

In 2010 Nashville experienced a 150-year record flood that inundated most of the downtown area and caused approximately \$1.5 billion worth of damage.¹ Such an occurrence is rare by definition, and experiencing it firsthand caused some “issues.”

An IT department located in the area had performed disaster planning and testing in the past, but never *full* disaster planning. So when this actual disaster occurred, they were well prepared for a disaster lasting a few hours.

No data was lost. Systems stayed online. But as the hours grew longer, they realized that they’d have to refuel their generators. Uh-oh. They had never done that before. And without power in much of the building, the IT workers had to lug five-gallon cans of diesel fuel up four floors every four hours . . . for a couple of weeks.

Perhaps gym memberships for IT personnel should be a component of your disaster recovery planning and testing

The goals of recovery practice are to discover problems with the recovery plan, to provide on-the-job training for key personnel to become familiar with the procedures and tools to be used for disaster recovery, and to raise awareness of the organization’s level of readiness to confront an actual disaster.

Of course, you might decide that it is impractical to conduct disaster recovery testing without advance warning. This is especially true if out-of-town travel is involved. However, to keep the test as close to the real thing as possible, do not use the advance warning as an opportunity to cheat (perhaps by sending additional materials to the off-site location that would not have been there otherwise).

The disaster recovery test should include all of the components of the written plan. This will include setting up the operating systems, installing the DBMS, recovering applications and data, and testing the recovered

environment for success or failure.

As mentioned earlier, you should periodically review the contents of your off-site tapes to ensure that they contain the correct backup data, and not just during a test. At the same time, you should review all additional materials sent off-site to assure that everything that is supposed to be there actually is there.

Personnel

Choosing the right team to design and carry out your disaster recovery plan is essential to the success of that plan. From the perspective of the DBMS, the disaster recovery team must be capable of installing and configuring the DBMS system software, assuring the integration of the DBMS with other system software components, recovering individual databases, testing the integrity of the databases, recovering related data that may not be stored in a database, installing and configuring application software, testing the applications, and taking care of the numerous details along the way. In other words, the team must be multiskilled and very adaptable. Moreover, the team members must possess a combination of IT skills and business skills.

Choosing the right team is essential.

Once the team is assembled, it is imperative that each member be trained to understand the ramifications of disaster recovery. This also means that the DBAs, programmers, and SAs on the team need to understand the business aspect of disaster recovery, and the business users need to understand the IT dimension of disaster recovery—at least at a very high level.

The entire team should meet regularly for cross training and perform a disaster recovery test at least annually. Remember that team members must be granted the proper authorizations at the remote site to perform their part of the disaster recovery. Nothing can stop a disaster recovery faster than lack of authority to perform a crucial task—especially if no one is around with the authority to grant authorizations, which could be the case during a disaster.

Backing Up the Database for Disaster Recovery

Your disaster recovery procedures will be determined in large part by the method you use to back up your data. If you rely on disk volume backups, your recovery will be one volume at a time. If you create image copies, you

will probably use the DBMS's RECOVER utility or a third-party recovery tool. Of course, you might combine several different techniques for off-site backups, depending on the sensitivity and criticality of the data. Let's look at several different backup strategies for disaster recovery.

Tape Backups

A valid strategy for disaster recovery backups is to use the same techniques deployed to create local backup files. Create multiple output files from the image copy backup process, and send tapes (or optical discs) of at least one of the copies to the remote disaster recovery site. Be sure to send the backup files to the disaster site as soon as possible after their creation. The sooner the backup media are shipped off-site, the less likely it is that they can be destroyed in a disaster.

You can use the same techniques deployed to create local backup files.

For disaster recovery purposes, you usually don't need to back up indexes. They can always be recreated from the data after it is recovered at the disaster recovery site.

Be sure to produce a daily report of all the backup files that were sent to the remote site. Keep a copy of the report at the local site and ship one to the remote site as well. The report should detail all of the information required to perform a recovery, including filename, type of backup (full or incremental), how the image copy was created (database utility or file system command), date and time of the backup, date and time it was shipped off-site, and the date and time it was received off-site (if possible).

Additionally, you will need to back up the database logs and ship them to the remote site. Failure to ship the log files to the remote site means that you will be able to recover each database object only to the point when the last backup was taken. Refer to [Figure 17.1](#). In this case, we have backed up the database object and sent a copy to the remote site. The copy may be sent to the remote site by packaging it up and physically shipping it to that location, or it may be shipped over a broadband connection to the remote location. Subsequent modifications occur to the database object and then a disaster strikes. Without a backup of the database log for the period between the image copy backup and the disaster, all of those changes will be lost. The

amount of data lost in an off-site recovery depends on the time-liness of the backup of database log files.

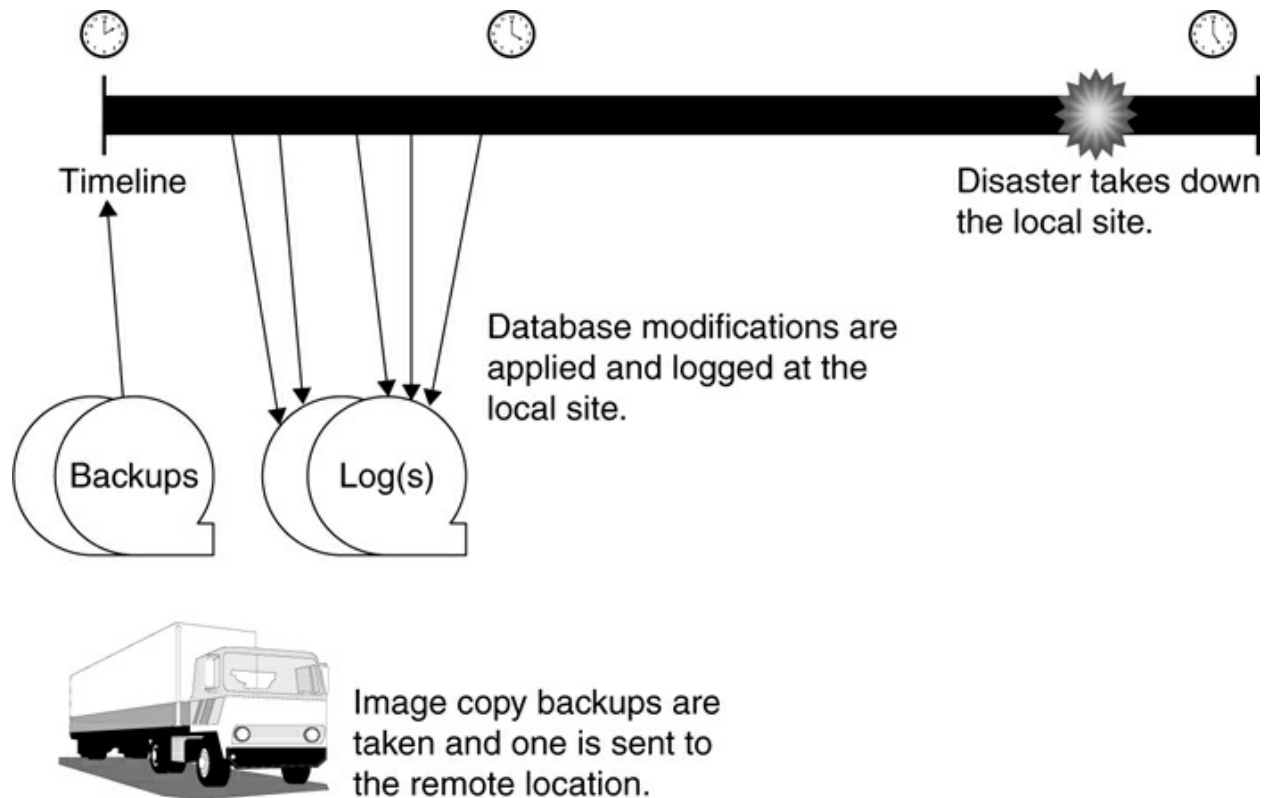


Figure 17.1. The database log and disasters

Back up the database logs and ship them to the remote site.

Recovery at the remote site is performed for each database object one by one. Indexes are rebuilt or recreated after the data is recovered. After recovery, the DBA should use whatever tools and utilities are available to identify and resolve any data integrity problems and constraint violations.

Additional preparation may be required, depending on the DBMS and operating system(s) in use. For example, if the DBMS keeps a record of the image copy backups, the system catalog will need to be recovered to a point where all of the off-site backups are listed in the system catalog—or some of the off-site image copy backups may not be able to be applied.

Of course, the system catalog will need to be recovered at the remote site regardless of the DBMS in order to recover any database objects at all. Other DBMS-related files may need to be recovered after the DBMS is installed as well.

Keep at least three backup tapes at your remote site for each database object. This provides a cushion in case one or more of the image copy tapes is damaged. The extra tapes will allow you to fall back to an older backup in the event of a tape failure. Furthermore, if you have the database logs, you may be able to recover the database object completely, even if the recovery does take longer.

Finally, when using this approach, be sure to consult the documentation provided by the DBMS vendor for the particulars of remote site recovery for that DBMS.

Consult the DBMS documentation for the particulars of remote site recovery.

Storage Management Backups

Another approach to disaster recovery backups is to use storage management software to make point-in-time copies of entire disk packs. Such an approach greatly simplifies disaster recovery preparation and execution, but this strategy can require a significant system outage to accomplish properly. To perform a backup using storage management software, follow these steps:

Use storage management software to make point-in-time copies of entire disk packs.

1. Stop the DBMS to create a system-wide point of stability for recovery.
2. Copy all of the database objects, using storage management software to dump complete disk volumes to tape.
3. When all of the disk volumes containing database objects have been successfully copied, restart the DBMS.
4. Copy the backup tapes and send them to the remote site.

Recovery at the remote site is then performed one complete disk volume at a time using the storage management software. The biggest problem with this approach is the requirement to stop the DBMS. Most organizations cannot tolerate such an outage due to e-business or global 24/7 requirements.

As an alternative to stopping the entire DBMS, data could be copied application by application. Using such an approach would require a regularly scheduled job to

1. Stop the application
2. Stop or quiesce the associated database objects
3. Back up the database objects with the storage management software
4. Restart the application and database objects after the backup is complete

Of course, such a compromise complicates the disaster planning and remote site recovery. The files containing the data for the application are not likely to exist on a single disk drive. Thus, identifying the correct disks and files to back up can be more difficult.

In general, stopping the DBMS and using storage management software for disaster recovery backups is effective for organizations that are willing to trade off around-the-clock availability for ease of disaster recovery preparation.

Other Approaches

Although the first two approaches are the most commonly implemented disaster recovery strategies, many other approaches and options exist. Deploying a wide-area network (WAN) to assist in the delivery of backups to the remote site is a good tactic to consider. If your primary site can be connected to the remote site using a WAN connection, you can direct your off-site backups to an electronic tape vault at the remote location. Of course, any direct connection between the primary site and the remote site increases the chances of a disaster at one location impacting the other location—although a network connection is probably a very small risk if implemented properly.

A WAN can be used to deliver backups to the remote site.

Another approach to disaster recovery preparation is the remote mirroring of data to the alternate site over the network. This approach minimizes the amount of preparation required for recovery at the remote site because the data is mirrored there. For this strategy to be effective, *all* changes at the primary site must be mirrored at the remote site, including regular database changes (INSERTs, UPDATEs, and DELETEs), database utility operations (LOAD and REORG), and local database recovery operations.

The standby database option for Oracle databases can be used as a disaster

recovery approach. As long as the standby database is located at a different physical site from the primary database, the standby database can be a good choice for implementing a disaster recovery plan for an Oracle instance.

Some Guidelines

In order for your recovery plan to be effective, you need to adhere to the written plan. Follow all of the routine precautionary measures documented in the plan to assure the recoverability of your critical data in the event of a disaster. Image copy backups or disk backups need to be made as directed and sent to the remote site as quickly as possible. Reports need to be printed and sent off-site. Missing any little detail can render a disaster recovery plan ineffective.

Adhere to the written plan.

When testing the disaster recovery plan, make sure that each team member follows the written instructions precisely. Of course, it is quite likely that things will come up during the tests that were missed or undocumented in the plan. Be sure to capture all of these events and update the written plan. Keep in mind that during an actual disaster you may need to rely on less experienced people, or perhaps consultants and others who are not regular employees. The more foolproof the written plan, the better the chance for a successful disaster recovery.

The following are some quick tips and hints to keep in mind as you develop your database disaster recovery plans.

Order of Recovery

Make sure the operating system and DBMS are installed at the correct version and maintenance level before proceeding with any database object recovery at the disaster site. Be sure to follow rigorously the recovery steps as documented in the written plan.

Data Latency

How old is the data? If you take nightly backup tapes to another location, your data could be up to 24 hours old. Sometimes having data that old is unacceptable, and sending backup media to off-site storage more than once a day is too expensive. One solution is to get the data to another location—via log shipping or replication, for example.

Some data may not be fully recoverable.

Database logs at the time of the disaster may not be available to apply at the off-site recovery location. Some data may not be fully recoverable, and there is really no way around this. The more quickly that backup copies of database objects and database logs are sent off-site, the better the disaster recovery will be in terms of data currency.

Remember Other Vital Data

Creating off-site backups for database objects may not be sufficient to ensure a complete disaster recovery plan for each application. Be sure to back up related data and send it off-site as well. Additional data and files to consider backing up for the remote site include

- DDL libraries for database objects, recovery and test scripts
- Application program source and executable files
- Stored procedure program source and executable files
- User-defined function source and executable files
- Libraries and passwords for critical third-party DBA tools
- Related data files used by the application

Beware of Compression and Encryption

If your site uses tape-compression software, be sure that the remote recovery site uses the same tape-compression software. If it does not, the image copy backups will not be readable at the remote site. Turn off compression at the primary site for the disaster recovery backups if the remote site cannot read compressed tape files.

Be sure that the remote recovery site uses the same tape-compression software.

Similarly, if the backup copies are encrypted, be sure that the decryption key is available at the remote site, otherwise the backup copies will not be able to be read, rendering them useless.

Post-Recovery Image Copies

An essential part of the disaster recovery process is to create an image copy backup for each database object after it has been recovered at the remote site.

Doing this enables easier recoverability of the data should an error occur after processing begins at the remote site. Without the new image copy backups, the disaster recovery procedure would have to be performed again if an error occurs after remote site processing begins.

Disaster Prevention

DBAs and IT professionals, in general, create procedures and enforce policies. Many of these procedures and policies, such as a disaster recovery plan, are geared toward dealing with situations once they occur. Having such procedures and policies is wise. However, isn't it just as wise to establish procedures and policies to prevent problems in the first place? Although you cannot prevent an earthquake or flood, you can implement policies to help prevent man-made disasters. For example, use surge protectors to prevent power surges from destroying computing equipment, and have backup generators on standby in case of electrical outages.

Establish procedures and policies to prevent problems in the first place.

Another good idea is to document and disseminate procedures to end users, teaching them how to deal with error messages. You cannot expect every user to understand the impact of responding to every error message. Depending on the application, the wrong user response to an error message can result in data loss. Guidelines can help avoid errors—and man-made disasters.

Disaster and Contingency Planning Web Sites

There are a number of useful sites on the Web that cover disaster and contingency planning in detail. The following sites can afford you a broader view of disaster recovery planning than the database-focused approach we have discussed in this chapter:

- www.thebci.org
- www.globalcontinuity.com
- www.sungard.com
- www.disasterrecovery.com/

Summary

A key part of a DBA's job is developing a plan to mitigate damage in the event of a disaster. The primary goal is to bring the applications back online with as little data loss and interruption as possible. Many components are required to create a viable disaster recovery plan, but it is imperative that a comprehensive written plan be created and maintained. Although the DBA's scope of responsibility is primarily assuring database availability and data integrity, the DBA must be part of a multidiscipline team for disaster recovery. Regular planning, testing, and revising are crucial to the creation of a usable disaster recovery plan.

The DBA must be part of a multidiscipline team for disaster recovery.

Review

1. What is a disaster?
2. Describe two different approaches to database disaster backup and recovery.
3. What factors should be considered when determining the criticality of a database object for disaster recovery planning?
4. What role does the database log play in disaster recovery?
5. What are the considerations for choosing a viable disaster recovery site?
6. Discuss the capabilities disaster recovery team members must possess to assure a viable recovery.
7. Once the disaster recovery plan is written it should not be changed: true or false?
8. Why is it a good idea to rank your applications and database objects for disaster recovery planning purposes?
9. With regard to data, what are the three categories of risk?
10. It is a good idea to perform a regularly planned disaster recovery test. However, what other events might cause you to test your disaster recovery plan more frequently?

Suggested Reading

- Bell, Judy Kay. *Disaster Survival Planning: A Practical Guide for Businesses*. Port Hueneme, CA: Disaster Survival Planning, Inc. (1991). ISBN 0-9630580-2-9.
- Bruni, Paulo, et al. *Disaster Recovery with DB2 UDB for z/OS*. Armonk, NY: IBM Redbook (2004). ISBN 0-7384-9092-X.
- Chantico Publishing Company. *Disaster Recovery Handbook*. Blue Ridge Summit, PA: Tab Professional & Reference Books (1991). ISBN 0-8306-7663-5.
- Chen, Whei-Jen, et al. *High Availability and Disaster Recovery Options for DB2 on Linux, UNIX, and Windows*. Armonk, NY: IBM Redbook (2009). ISBN 0-7384-3138-9.
- Mullins, Craig S. *DB2 Developer's Guide*. 6th ed. Boston, MA: IBM Press (2012). ISBN 978-0-13-283642-5.
- Photopoulos, Constantine. *Managing Catastrophic Loss of Sensitive Data*. Burlington, MA: Syngress (2008). ISBN 978-1-59749-239-3.
- SunGard Recovery Services, Inc. *Action Plan for Disaster* (1995).
- Toigo, Jon William. *Disaster Recovery Planning*. 2nd ed. Upper Saddle River, NJ: Prentice Hall (2000). ISBN 0-13-084506-X.

18. Data and Storage Management

All DBMSs rely on data files to store data, and these files, or *data sets*, reside on storage media, or devices. Storage management is thus a key part of the database operations required of a DBA. Typically, storage means fixed disk drives or disk subsystems, but storage can also mean nonvolatile storage, solid-state devices, removable storage, optical storage, or tape devices.

Since the DBA deals with database management systems that store data, a good part of the DBA's job will involve planning for the actual storage of database data. In order to perform this part of the job, the DBA must be conversant in the actual physical mechanisms available for data storage. Furthermore, the DBA must understand the ways in which the abstract concept of data as embodied by the database interacts with the physical storage of data on persistent storage media.

Storage Management Basics

As a rule, DBMS vendors do not certify or explicitly support any specific third-party storage products. Instead, the assumption is made that some underlying storage technology is available and will be reliable. To determine what will work best with each DBMS, the DBA must evaluate the many products, technologies, and vendors that provide storage solutions.¹ Although most storage technologies work with most DBMS products, some storage technologies are better suited than others in terms of performance, reliability, usability, and cost.

Some storage technologies are better suited than others.

The predominant storage technology used for data management is the disk drive. Modern disk drives are more reliable than in years past, with an ever-increasing mean time between failures (MTBF). It is not unheard of for disk drives to achieve in excess of 100,000 hours of availability before failing. However, the mechanical nature of disk drives renders them more vulnerable to failure than other computer components. As the number of disk drives in a system increases, the vulnerability of the system increases. It is not uncommon for a single organization to rely on hundreds or thousands of disk

drives to support its database applications. Furthermore, the ways in which databases use disk drives can make storage management unpredictable, as the sidebar “[Modern DBMS Disk Usage](#)” points out. Certain modern storage solutions, such as RAID, can be used to address some of the MTBF problems.

However, for mission-critical applications data *integrity* can be more important than data *availability*. If the storage media is unreliable and a failure causes data corruption, the lost data can be more of a problem than the downtime. It is imperative, therefore, that database storage solutions protect the data at all costs.

Modern DBMS Disk Usage

DBAs must be prepared for spikes in disk usage. The inherent functionality of modern DBMSs makes them much more likely than their forebears to cause disk usage spikes. Some DBMSs allow the temporary file sizes to expand and contract automatically. Depending on the type and nature of database operations in process, this fluctuation can cause disk usage spikes. Additionally, many online utility operations use shadow objects. Therefore, if a tablespace normally occupies 10GB of disk space, an online reorg would require 20GB—10GB for the normal object and another 10GB for the shadow object.

Database performance is I/O dependent—the faster the DBMS can complete an I/O operation, the faster the database application will run. Remember that data retrieval from storage media takes much longer to complete than data retrieval from cache or memory. For this reason, some modern storage systems provide their own caching mechanism to prestage data in memory—thereby reducing the wait time associated with traditional disk I/O operations.

Database performance is I/O dependent.

Indeed, storage is becoming more central to business operations. Heterogeneous, multiterabyte database sites are not uncommon these days. The amount of data being stored is greater than ever before, and the data is being stored for longer durations. Lou Gerstner,² former CEO of IBM Corporation, summed up the phenomenal rate of data growth: “Inside IBM

we talk about ten times more connected people, a hundred times more network speed, a thousand times more devices, and a million times more data.” Mr. Gerstner is not the only one to note this phenomenal growth:

- According to analysts at IDC Corporation, the amount of digital information created in the world in 2010 exceeded a zettabyte (1 trillion gigabytes) for the first time.³ IDC went on to predict that the amount of digital information created in 2011 would surpass 1.8 zettabytes.
- A 2011 study⁴ offers additional insight into the rapid growth rate of data storage. According to the study, almost all shops reported data growth over the previous year, and one-third reported the amount of data within their enterprises grew by 25 percent or more during the prior year. Additionally, about 10 percent had data stores in the petabyte range. Furthermore, 27 percent of those surveyed reported environments exceeding 100TB, whereas only 20 percent had reported that size in the previous survey. The net result is that organizations are storing more and more data.
- Winter Corporation publishes a regular survey of the largest and most heavily used databases in operation, called the TopTen Program.⁵ The most recent report highlights a 23TB transaction processing database hosted on the mainframe and a 100TB data warehouse hosted on UNIX.
- A now somewhat dated research note⁶ from the Giga Research Group tackled the question “How much data is on the planet?” That research estimated the amount of data to be approximately 201,000TB, or about 197PB. Of course, this is just an estimate that Giga believes to be accurate within an order of magnitude, that is, a factor of ten.

Indeed, talking about terabyte- and petabyte-size databases and data warehouses is becoming common these days. [Table 18.1](#) outlines the measurements used when discussing data storage size. Keep in mind, though, that the figures in this chart are rough guides. Some people speak about disk storage in terms of “powers of ten” instead of “powers of two”—in other words, they refer to a kilobyte as 1,000 bytes instead of 1,024 bytes. However you choose to measure it, though, disk drive capacity is increasing very rapidly.

Table 18.1. Data Storage and Size Terminology

Abbreviation	Term	Size	Power of 2
B	Byte	8 bits	
KB	Kilobyte	1,024 bytes	2^{10} bytes
MB	Megabyte	1,024KB	2^{20} bytes
GB	Gigabyte	1,024MB	2^{30} bytes
TB	Terabyte	1,024GB	2^{40} bytes
PB	Petabyte	1,024TB	2^{50} bytes
EB	Exabyte	1,024PB	2^{60} bytes
ZB	Zettabyte	1,024EB	2^{70} bytes
YB	Yottabyte	1,024ZB	2^{80} bytes

The growth in storage capacity further increases the complexity of managing data and databases. Many organizations are implementing new storage technologies, such as network-attached storage (NAS) and storage area networks (SANs), to help manage the ever-increasing amount of storage required for modern applications. Managing storage in today's dynamic environment is a challenging DBA task.

Managing storage in today's dynamic environment is a challenging task.

Goals to consider while building a storage system include

- Preventing loss of data—the number-one priority
- Assuring that adequate capacity is available and that the storage solution can easily scale as storage needs grow
- Selecting a solution that provides fast access to data with minimal or no interruptions to service
- Choosing storage solutions that are fault tolerant and that can be repaired quickly when a failure occurs
- Selecting a storage solution that allows you to add or replace disks without an outage
- Combining all of the above into a cost-effective storage solution your company can afford

Files and Data Sets

There are many storage issues that must be resolved before a DBA can create a database. One of the most important issues is how much space to allow for the database. A data file corresponds to the tablespace database object as shown in [Figure 18.1](#). Keep in mind that the space calculation must take into account not just tablespaces, but also indexes and, depending on the DBMS, the transaction log. Each of these entities will likely require a separate file, or data set, for persistent storage.




Figure 18.1. Database storage and data files

After determining the amount of storage required, the DBA must choose an appropriate storage device with sufficient space to store the database and all of its files. The DBA may choose to use multiple storage devices for the different files to

- Align the performance requirements of the file with the appropriate disk device
- Separate indexes from data for performance reasons
- Isolate the transaction log on a separate and very fast device
- Isolate temporary and work files on a single volume; if a disk error occurs, temporary files can be deleted and redefined with no backup and recovery implications

- Spread the data across multiple devices to facilitate parallel access

The DBA may choose to use multiple storage devices.

File Placement on Disk

The DBA must determine the optimal placement of files on disk devices. At times the DBA can achieve performance gains simply by moving files from one physical disk device to another. One common technique is to place index files and data files on separate disk devices. When indexes are separated from the data that is indexed, I/O operations can be made more efficient because the physical read-write arm of a single disk device does not need to move multiple times. For the same reason, placing data that is accessed by the same operations on separate physical disk devices is another common file placement technique and provides the same type of performance advantage as separating indexes from data.

Short-stroking disks is another technique that can be used to maximize the performance of disk subsystems. The general idea is to only use the very outside edge of the disks so that the arm does not need to move very far on the platter to find the needed data. Short-stroke disk techniques also keep the platter from having to rotate very many times to find all that data. The outer ring of the platter holds more data per rotation than the inner rings of the platter just due to the simple fact that there is more surface area.

Short-stroke disk techniques keep the platter from having to rotate very many times to find all that data.

Of course, exact file placement is becoming less beneficial with the advent of modern storage devices. If the DBMS is using a modern storage device that creates a virtual disk by spreading the data across multiple physical disks (RAID), explicit file placement is a waste of time. Even if the DBA specifies two different disk devices, the disk array will physically place the files across multiple disks. So, as a DBA, do not waste time on precise file placement when using a modern storage array technology such as RAID.

Regardless of the type of storage being used, be sure to place the transaction log on a separate device from the database so that the transaction log can be backed up independently from the database. This is a good idea because it increases overall recoverability.

Place the transaction log on a separate device from the database.

Every DBMS provides different storage options. Microsoft SQL Server offers filegroups, DB2 for z/OS provides STOGROUPS, and Sybase offers segments (see the sidebar). Be sure to understand the mechanisms your DBMS uses to interact with the storage subsystems and disks to create database files. Improperly created database files can be a significant cause of poor performance.

Some organizations choose to implement system-managed storage, or SMS. With SMS, the actual location of files and data sets is determined by the system, instead of a DBA or storage administrator. With the new efficient disk technology that is available, SMS is a more viable option than it was in the past.

SMS is a more viable option than it was in the past.

Sybase Segments

Some databases provide options for controlling the size of tables. For example, Sybase segments can be used to limit table size because a table can never grow larger than its segment allocation. To control growth of a table, the DBA can create it on a specific segment of a specific size. The table will never be able to grow larger (without taking explicit action to extend the segment to another device).

Indexes and tables can be allocated to specific segments to enhance performance. Placing nonclustered indexes on separate segments from the table reduces disk head contention if the segments reside on separate physical disks.

An ideal use of segments is to isolate large, frequently accessed tables. In addition, segments can be used to split a single table across two disks.

One possible scenario is to let SMS handle the placement of the majority of your database files. However, for certain performance-sensitive files (such as the transaction log or mission-critical database objects), use non-SMS file placement techniques. With this approach, SMS can minimize the effort

involved in most database file placement tasks, and the DBA can get involved only for “high-need” database file placement.

Raw Partitions versus File Systems

A raw partition is the preferred type of physical device to choose when deploying databases on UNIX-based database servers. A *raw partition* is simply a raw disk device with no operating system or file system installed.

Raw partitions are favored because of the way UNIX buffers writes when a file system is used. When writes are buffered by the operating system, the DBMS cannot know whether the data has been physically copied to disk or not. When the DBMS writes data from cache to disk, UNIX in all probability will not. If a failure occurs, data in a database using a disk system file may not be 100 percent recoverable. This should be avoided.

If a raw partition is used, the data will be physically written to disk with no operating system intervention. Additionally, when using a raw partition, the DBMS is better able to ensure that enough space is available and write the allocation pages as needed. When using a disk file, UNIX will not preallocate the space required; it will instead create a sparse file. Only the allocation pages are written. If other applications are contending for the same physical space, you may not have as much space for storing database data as you believe you do.

The drawback to raw devices is the difficulty of tracking the database files. Since the files were not allocated using the operating system or file system, it is not possible to track them using operating system and file system commands. However, third-party storage management software can be used to get around these problems.

The drawback to raw devices is the difficulty of tracking the database files.

On UNIX, be sure to favor using a raw partition for transaction logs, even if you use a file system for database files. It is usually wise to use raw partitions for production databases, as well. However, if integrity is not an issue (for example, in a development environment), databases can be created on a disk file system for the sake of simplicity.

Temporary Database Files

Modern DBMSs provide capabilities to create temporary database objects that exist only during the scope of a specific transaction. These objects contain temporary data that is transient and does not require long-term persistent storage. However, these database objects still require some form of short-term persistent storage to use for the duration of their existence.

Temporary database objects require some form of short-term persistent storage.

Depending on the DBMS, the DBA will need to assign disk devices and an amount of storage for use by temporary database objects.

Space Management

As modifications are applied to database tables, databases will grow in size. Remember, too, that a database is not just the data portion (tables and indexes), but also the log portion. It is wise to periodically and consistently monitor database space usage. This can be done using the tools and utilities provided with the DBMS, storage management software, or third-party database tools. As a DBA, you should be able to track the following:

- Number of secondary extents
- Device fragmentation
- Fragment usage information
- Free space available
- Segment or partition size
- Tables and indexes allocated per segment
- Amount of reserved space that is currently unused
- Objects approaching an “out of space” condition

Consistently monitor database space usage.

Dealing with a database that has run out of space is a common nuisance for DBAs. But it will occur, so it is best to be prepared with a plan of attack for handling database growth. The first step is to reclaim space by regularly dropping unused database objects. Objects are sometimes created, forgotten, and never required again. When unused database objects are dropped, the space they previously occupied can be made available.

If no unused database objects exist, the DBA must expand the storage available for the database. Typically, this is accomplished using the ALTER statement to specify additional space for the database object. The DBA also may have to allocate an additional storage device to be used by the database object, which may involve moving the files for the database object from one device to another. Be careful when specifying additional space for database objects using the ALTER statement. Some DBMSs require the DBA to specify the amount of increase, whereas others require the specification of a new overall size.

Data Page Layouts

Each DBMS uses a different format for laying out the pages of database objects. Typically, though, the page layout for a database object will consist of three basic components, as shown in [Figure 18.2](#).

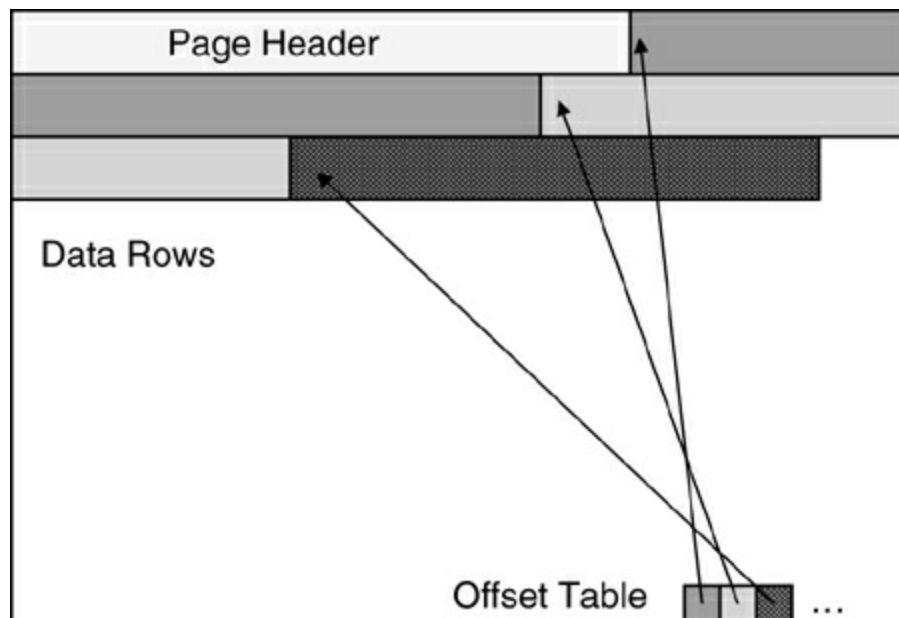


Figure 18.2. Sample data page layout

The page for a database object usually consists of three components.

- *Page header.* The very beginning section of a database page will contain several bytes of physical “housekeeping” information. The page header may include a page identifier, forward and backward links to other data pages, an identifier indicating to which table the page

belongs, free space pointers, and the minimum row length for the table.

- *Data rows.* A page also will contain the actual rows of the table (or index) containing the user data. Rows will not cross a page boundary except for certain large data types such as text, images, and other binary large objects.
- *Offset table.* An offset table may exist with pointers to each data row on the data page. Some DBMSs always use some form of offset table, whereas others use the offset table only when variable-length rows exist on the data page.

Keep in mind that the actual layout and length of the components of a page will differ from DBMS to DBMS. In addition, the number of rows that can be stored on a database page will depend on the page size. Some DBMSs require a fixed page size for database objects, whereas others permit the DBA to specify the page size.

Allocation Pages

In addition to data pages, the DBMS will use another type of page to control space usage. Sometimes called a *space map page*, an *allocation page* is used by the DBMS to manage the other pages in the database object. Multiple allocation pages can exist within a database, each one controlling a block of other data pages. These blocks are typically referred to as *allocation units*.

The DBMS uses an allocation page to manage the other pages in the database object.

Allocation pages control physical pages. Each physical page is mapped to a single database. Physical pages need not be contiguous on disk and, in fact, frequently will not be. The primary vehicle used by an allocation unit to manage page allocation is the bitmap. A bitmap contains a series of bits that are either turned on (“1”) or off (“0”). Each bit refers to pages within the allocation unit and to whether space is available within the page. The DBMS uses this information to facilitate quicker data modification processes. The DBMS may use multiple bitmaps to manage allocation, deallocation, and space usage. The allocation page will have a distinctly different format from a typical data page.

The DBA must factor the existence of allocation pages into any space calculations performed to manage the storage used by the database. The DBA

will need to know the number of pages in each allocation unit that is controlled by each allocation page. The DBA then adds this additional overhead into the database space usage calculations.

Data Record Layouts

A table is composed of rows and columns. Each single row must be entirely stored on a single data page (with the exception of some large data objects in some DBMSs). Each row is contained within a record that consists of additional housekeeping data along with the actual data contents of the row. Each table record comprises the following elements:

- *Row header.* Each record typically begins with several bytes of physical “housekeeping” information detailing the structure and composition of the data contents of the row. This might include row length, information on variable-length data, and other control structures.
- *Row data.* This consists of the actual data contents of the data columns for the row, in the order of their definition. Depending on the DBMS, the variable- and fixed-length columns may be separated.
- *Offset tables.* Optionally, the record may contain an offset table with pointers to manage and control where variable-length fields are stored within the row.

A record contains a row’s data contents as well as housekeeping data.

When calculating the storage required for a table, the DBA must calculate record length, not row length. Record length includes the header and offset table overhead, whereas row length is simply a count of the length of each column in the table. Of course, the DBA must be able to accurately calculate row length, too, because the row length is a component of the overall record length.

To calculate the length of the data portion of the row, the DBA will need to secure documentation that outlines the way to determine the actual, physical length for each data type supported by the DBMS. The actual data types and lengths required to store columns of these data types differ from DBMS to DBMS. The following formulas can be used as general guidelines for calculating the physical size of a row:

Fixed-length rows:

$$\begin{aligned} \text{Row Size} &= \text{Length (COL1)} \\ &+ \dots \\ &+ \text{Length (COLn)} \end{aligned}$$

Variable-length rows:

$$\begin{aligned} \text{Row Size} &= \text{Length (FIXED COL1)} \\ &+ \dots \\ &+ \text{Length (FIXED COLn)} \\ &+ \text{Average [Length(VAR COL1)]} \\ &+ \dots \\ &+ \text{Average [Length(VAR COLn)]} \end{aligned}$$

The length of a column will depend on the data type of that column; each DBMS documents the length of the data types it supports. To calculate the length of the entire record you will need to find the amount of overhead required for the type of rows (variable or fixed length) and add that to either calculation.

The length of a column will depend on the data type of that column.

Calculating Table Size

Once the page layouts are known and the DBA can accurately calculate record lengths, calculating the size of a table is easy. After calculating the record length, the next step is to determine how many rows will fit on a physical page. The size of a page will differ from DBMS to DBMS, and most DBMSs allow the DBA to specify different page sizes.

The size of a page differs from DBMS to DBMS.

For example, assume that a single page consists of 4K of storage, but 32 bytes of that page are required for overhead (header information). This leaves 4,064 remaining bytes in which data rows can be stored. Use the following formula to arrive at the number of rows that can fit on a single data page:

$$\text{Rows per page} = \frac{4064}{\text{Row size}}$$

Use the following formula to arrive at the total amount of space required for the table:

$$\text{Table size (in K)} = \frac{\text{Number of rows}}{\text{Rows per page}} \times (4K)$$

Of course, the DBA will need to factor in the free space that is specified for the table as well. Furthermore, remember to factor in data that is stored externally from the rest of the data, such as text columns in SQL Server or (most) LOB columns in DB2.

Index Page Layouts

When a table has no indexes defined for it, rows are simply added at the end of the table. When there is no additional room to insert new rows on the last page, a new page is obtained (either from the current allocation or by creating a new extent). The data pages of a nonindexed table are linked together with forward and backward pointers. Traversing the linked list causes the data pages to be accessed in order.

Similar restrictions apply to records stored in an index as apply to a data record. A single index row must be entirely stored on a single index page. The typical index record includes the following items:

- *Header information.* As with data pages, each index record typically begins with several bytes of physical “housekeeping” information detailing the structure and composition of the index record.
- *Row length.* For variable-length keys, the index may need to store the actual length of the indexed data.
- *Index key values.* The actual data values for the index key are listed in order of definition.
- *Page pointer.* This points to the physical location of the data page in the table that actually holds the indexed data.
- *Offset and adjust tables.* These may be required to manage and control the position of variable-length fields stored within the index row.

A single index row must be entirely stored on a single index page.

The composition of the pointer component of the index record depends on whether the index is clustered or nonclustered. Clustered indexes require only

a pointer to the data page. Nonclustered indexes require both the pointer to the data page and a pointer to the location of the data row within the data page.

Calculating Index Size

The first step in calculating the space required to store an index is to calculate the row size for the index using one of the following formulas:

Fixed-length index keys:

$$\begin{aligned} \text{Row Size} &= \text{Length (KEY-COL1)} \\ &+ \dots \\ &+ \text{Length (KEY-COLn)} \end{aligned}$$

Variable-length index keys:

$$\begin{aligned} \text{Row Size} &= \text{Length (FIXED-KEY-COL1)} \\ &+ \dots \\ &+ \text{Length (FIXED-KEY-COLn)} \\ &+ \text{Average [Length (VAR-KEY-COL1)]} \\ &+ \dots \\ &+ \text{Average [Length (VAR-KEY-COLn)]} \end{aligned}$$

Remember, you need to calculate the size of an index record, not just the size of the row. Records are simply the row plus the overhead required to store the row. To calculate the record size you will need to obtain the size of the row overhead required for indexes in the DBMS you are using.

You need to calculate the size of an index record, not just the size of the row.

Once you have calculated the record size, you can move on to the next step, which is to calculate the amount of storage space required for the index. Begin by using the following formula to arrive at the number of records that can fit on a single index page:

$$\begin{aligned} \text{Entries per page} &= (\text{Block size} - \text{Overhead}) \\ &\div \text{Row size} \end{aligned}$$

Then you can use the following formulas to arrive at the total amount of space required for the index:

$$\begin{aligned}
 \text{First-level pages} &= \text{Number of rows} \\
 &\div \text{Entries per page} \\
 \\
 \text{First-level size} &= \text{First-level pages} \\
 &\times (\text{Block size} - \text{Page overhead}) \\
 \\
 \text{Second-level pages} &= \text{First-level pages} \\
 &\div \text{Entries per page} \\
 \\
 \text{Second-level size} &= \text{Second-level pages} \\
 &\times (\text{Block size} - \text{Page overhead}) \\
 \\
 \text{Third-level pages} &= \text{Second-level pages} \\
 &\div \text{Entries per page} \\
 \\
 \text{Third-level size} &= \text{Third-level pages} \\
 &\times (\text{Block size} - \text{Page overhead})
 \end{aligned}$$

And so on. Continue the iteration until the number of pages is one or fewer. At that point you have reached the root level of the index. Adding together the sizes for each level will result in the total amount of space required to store the index.

Transaction Logs

One of the most important files for database processing is the transaction log. Remember, every insert, update, and delete to any database table is recorded (or logged) in the transaction log. Arriving at an exact size for the transaction log can be more of an art than a science, because transaction activity is by nature random.

Arriving at an exact size for the transaction log can be more of an art than a science.

Database activity is the most important factor in determining an appropriate size for the transaction log. A good rule of thumb is to provide sufficient space for the log file to capture all database modifications that will occur between log archivals during your busiest processing period. The tricky

part is determining the most active processing period and the database modification's impact on the log. At any rate, if you do not automatically archive or back up your database logs, you should pay close attention to the size of the database log file(s). Failure to do so can cause data to be unrecoverable or cause processing to slow down or halt completely.

The DBA must specify other database logging parameters that impact disk usage. Most DBMSs support multiple, redundant logs. The DBA must decide how many transaction logs there should be and on which disk devices they should be allocated. At a minimum, you should run with dual active logging—that way, a problem on one log will not shut down DBMS operations, because the other active log can be used.

Another consideration is how long archive logs should remain on disk. For some implementations, logs can be archived directly to disk. However, from an efficiency standpoint, it makes more sense to archive logs to disk, and then have them roll off to tape later. As a rule of thumb, consider keeping your archive logs on disk for at least 24 hours. Furthermore, plan ahead to make sure you have sufficient disk space to accommodate archive logs for the system throughput on your busiest day.

Fragmentation and Storage

Fragmentation can be the enemy of performance and storage requirements. But let's not get ahead of ourselves: What is fragmentation? The common definition is the scattering of parts of a file throughout a disk, such as when the operating system breaks up a file and fits it into the spaces left vacant by previously deleted files. There are several aspects of fragmentation that the DBA must monitor for and correct.

Fragmentation is the enemy of performance and storage requirements.

Disk fragmentation can impact performance on some systems. Disk fragmentation occurs at the operating system level, and the DBA must be able to use OS tools to analyze drives on which databases reside. On small-scale environments with conventional I/O subsystems, it is a good idea to correct disk fragmentation before running defragmentation tools provided by the DBMS. On large-scale environments that benefit from more intelligent disk subsystems, correcting disk fragmentation may not be needed.

Index fragmentation is a more serious concern for DBAs. As discussed in more depth in [Chapter 11](#), indexes can become disorganized and fragmented as data is added, modified, and removed from the tables upon which the index is built. Fragmented indexes consist of many scattered areas of storage that are too small to be used productively. This causes wasted space, which can hinder performance and increase storage costs. Other conditions that can impact storage usage include declustered data, page splits, and row chaining.

Fragmentation impacts performance because not only is the data suboptimally located, but when a database read engine scans indexes for read-ahead processing, fragmentation can cause additional data pages to be scanned into the data cache, thereby degrading performance.

DBAs should use tools⁷ to scan database indexes for fragmentation and take actions to defragment or rebuild indexes on a regular basis. Doing so will not only improve performance but also reduce the amount of storage required for indexes.

Storage Options

Once you have calculated the amount of storage required for your database files, you will need to choose the type of storage to be used. There are several storage options available from which the DBA can choose, based on fault tolerance, performance requirements, and budget.

The primary storage medium used by database files will, of course, be direct-access disk drives. But other types of media are available, such as solid-state devices (SSDs), also referred to as solid-state drives, solid-state disks, or electronic disks. An SSD uses solid-state memory to store persistent data in a manner similar to a hard disk drive. But because SSDs use memory as opposed to magnetic disks, which are mechanical devices with moving parts, I/O performance is greatly improved.

Although SSDs outperform traditional disks and are less susceptible to physical damage, they are considerably more expensive and tend to have lower capacities.

SSDs outperform traditional disks and are less susceptible to physical damage, but they are more expensive and have lower capacities.

Some DBMS products, called *main-memory database management systems* (MMDBMSs), are specifically engineered for the persistent storage of data in memory. The most popular of these, TimesTen and SolidDB, were acquired by Oracle (2005) and IBM (2008) respectively. Although the popularity of MMDBMS products was rising in the middle 2000s, it looks increasingly as if MMDBMS will remain a niche offering, as traditional DBMSs can be coupled with SSDs to produce similar results.

Tape is also a storage option, as is optical disc, but these are WORM (write once, read many) technologies. Tape is used primarily for backup purposes, while optical disc is a niche technology used primarily for backup and offline data storage.

Which brings us back to good old disks. Disk devices are the predominant storage vehicles used by DBAs for persistent data storage. Disk drives are designed for direct read-write access, are cost-efficient, and are capable of storing very large volumes of data. Some disk devices are slower than others; some are smaller. However, an astounding variety of disk drives are available in various speeds, capacities, and price ranges to satisfy just about any type of storage requirement.

Disk devices are the predominant storage vehicles used for persistent data storage.

RAID

Outages due to simple media failures often can be avoided by implementing modern disk technologies such as RAID. *RAID*, an acronym for *redundant arrays of independent disks*, combines multiple disk devices into an array that is perceived by the system as a single disk drive. There are many levels of RAID technology, which deliver different degrees of fault tolerance and performance.

RAID combines multiple disk devices into an array.

Another desirable aspect of RAID arrays is their ability to use hot-swappable drives; the array does not have to be powered down to replace a failed drive. That a drive can be replaced while the array is up and running is a plus for DBAs needing to manage databases while delivering nonstop data availability.

RAID Levels

Several different levels of RAID are available. Vendors provide varying levels of support for the RAID levels that have been defined. These various levels of RAID support continuous availability through combinations of functions called mirroring, striping, and parity. *Mirroring* occurs when complete copies of the data are made on at least two disk drives, and all changes made to the data are made simultaneously to both copies. If one fails, access is automatically shifted to the remaining copy. *Striping* occurs when subsets of data are spread across multiple disk drives. If any one drive fails, the impact of the failure is limited to the data within the stripe on that disk. Finally, *parity* bits are encoded data that can be used to facilitate the reconstruction of the original data, in the event that all or part of the data cannot be accessed if the drive fails. The lost data can be reconstructed “on the fly” until it can be rewritten to undamaged disks.

Vendors provide varying levels of support for RAID levels.

Raid-0 (see [Figure 18.3](#)) is also commonly referred to as *disk striping*. With RAID-0, data is split across multiple drives, which delivers higher data throughput. However, there is no redundancy (which really doesn't fit the definition of the RAID acronym). Because no redundant data is being stored, performance is usually very good, but a failure of any disk in the array will result in data loss. Thus, RAID-0 should be avoided in mission-critical implementations.

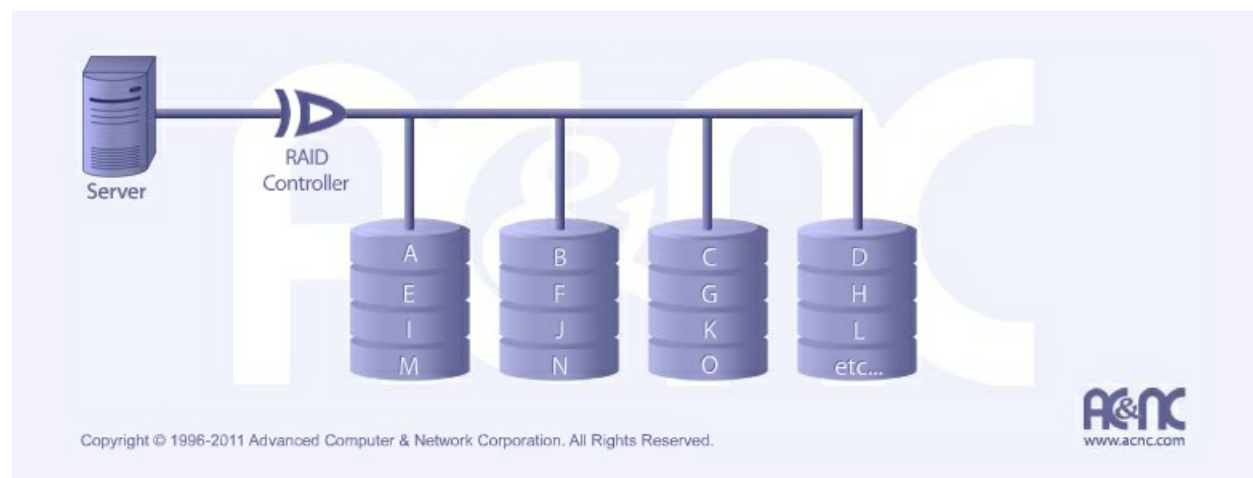


Figure courtesy AC&NC, www.raid.com.

Figure 18.3. RAID-0

RAID-1 (see [Figure 18.4](#)), sometimes referred to as *data mirroring*, provides redundancy because all data is written to two or more drives. A RAID-1 array will generally perform better than a single drive when reading data and worse when writing data. However, RAID-1 provides data redundancy, so if any drive fails, no data will be lost.

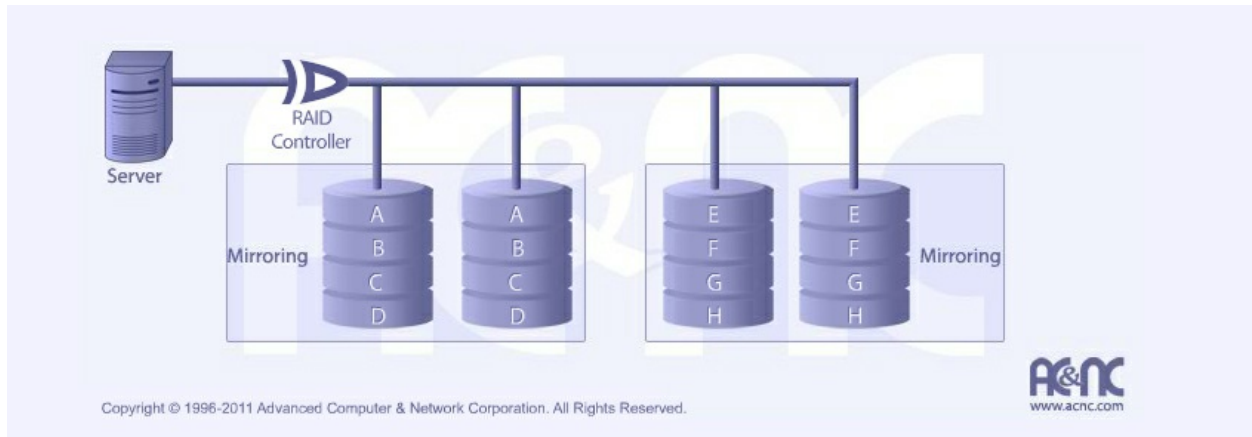


Figure courtesy AC&NC, www.raid.com.

Figure 18.4. RAID-1

RAID-2 (see [Figure 18.5](#)) provides error correction coding. RAID-2 is useful only for drives without any built-in error detection. No commercial RAID-2 implementations exist because it is not a commercially viable technique.

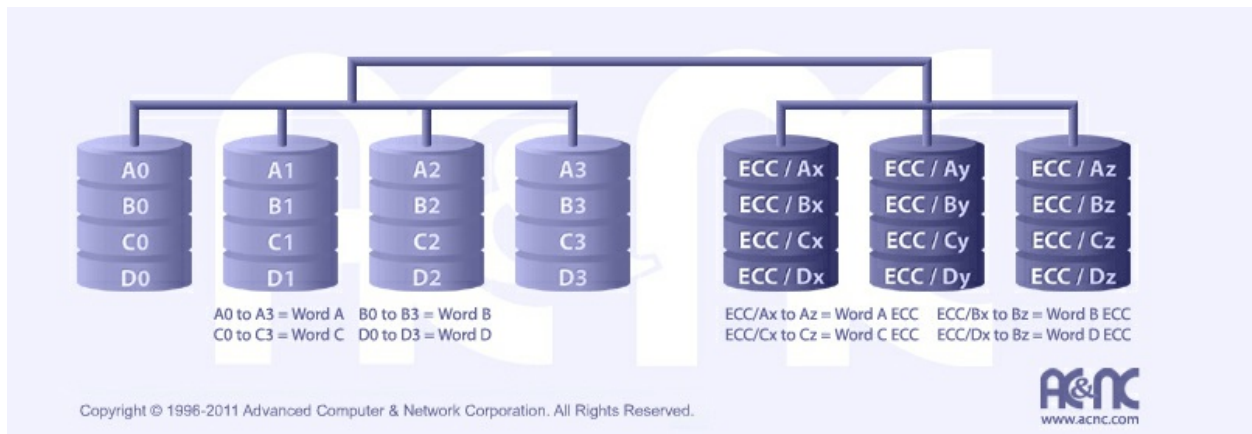


Figure courtesy AC&NC, www.raid.com.

Figure 18.5. RAID-2

RAID-3 (see [Figure 18.6](#)) stripes data at a byte level across several drives, with parity stored on one drive. RAID-3 provides very good data transfer rates for both reads and writes, but it requires a fairly complex controller to implement.

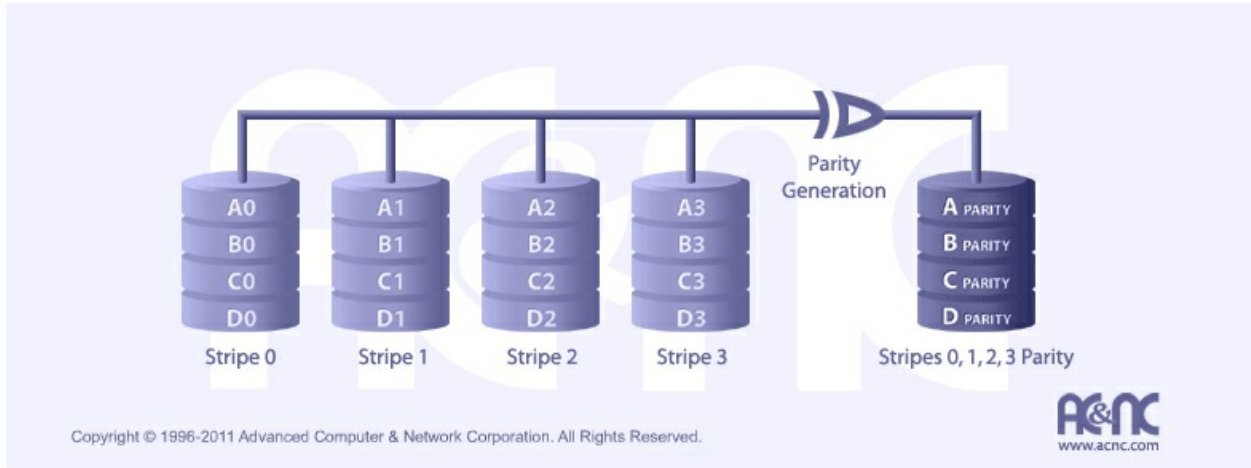


Figure courtesy AC&NC, www.raid.com.

Figure 18.6. RAID-3

RAID-4 (see [Figure 18.7](#)) stripes data at a block level across several drives, with parity stored on a single drive. For RAID-3 and RAID-4, the parity information allows recovery from the failure of any single drive. The write performance can be slow with RAID-4, and it can be quite difficult to rebuild data in the event of RAID-4 disk failure. RAID-4 also requires a complex controller design.

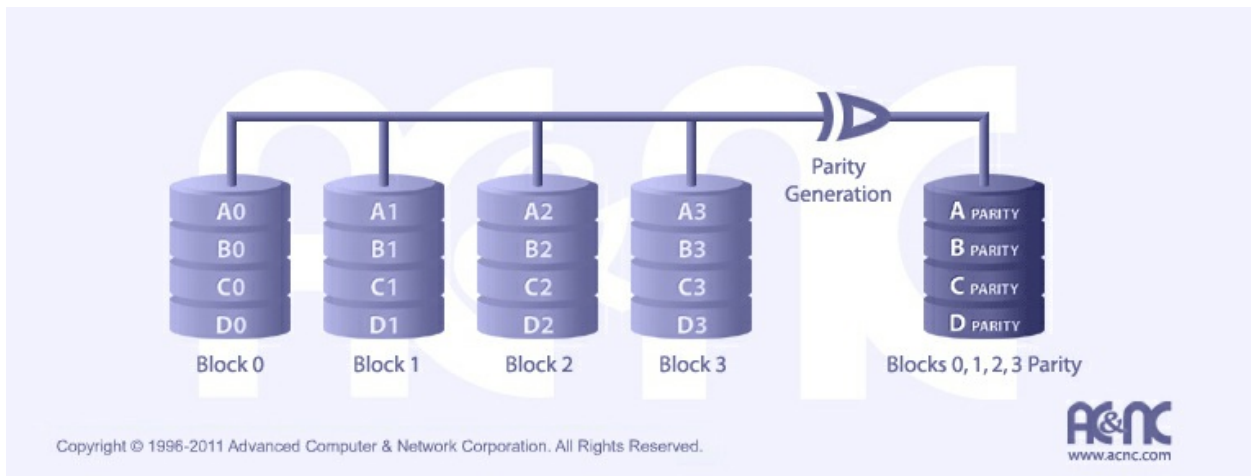


Figure courtesy AC&NC, www.raid.com.

Figure 18.7. RAID-4

RAID-5 (see [Figure 18.8](#)) is similar to RAID-4, but it distributes the parity information among the drives. RAID-5 can outperform RAID-4 for small writes in multiprocessing systems because the parity disk does not become a bottleneck. However, read performance can suffer because the parity information is on several disks.

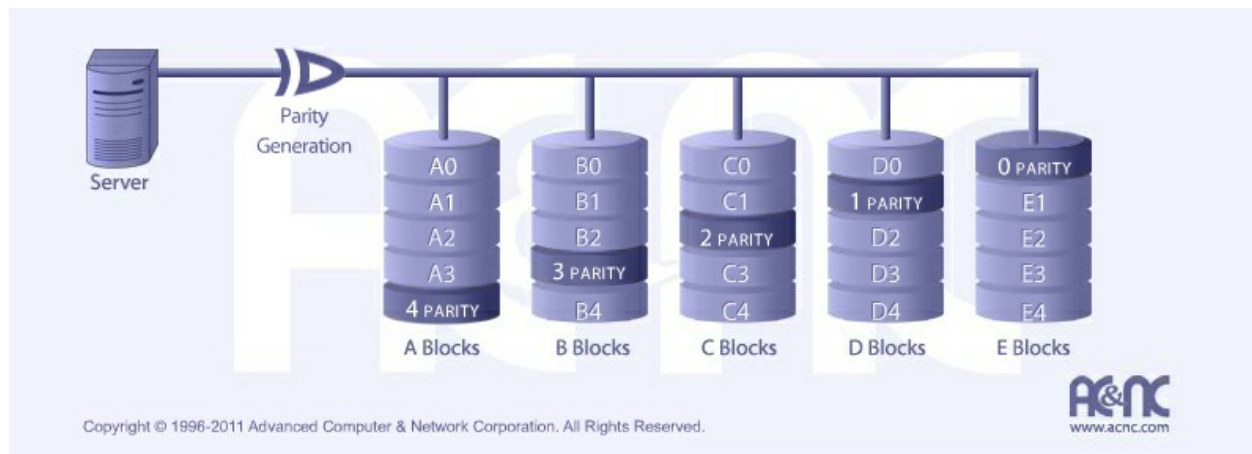


Figure courtesy AC&NC, www.raid.com.

Figure 18.8. RAID-5

RAID-6 (see [Figure 18.9](#)) is basically an extension of RAID-5, but it provides additional fault tolerance through the use of a second independent distributed parity scheme. Write performance of RAID-6 can be poor.

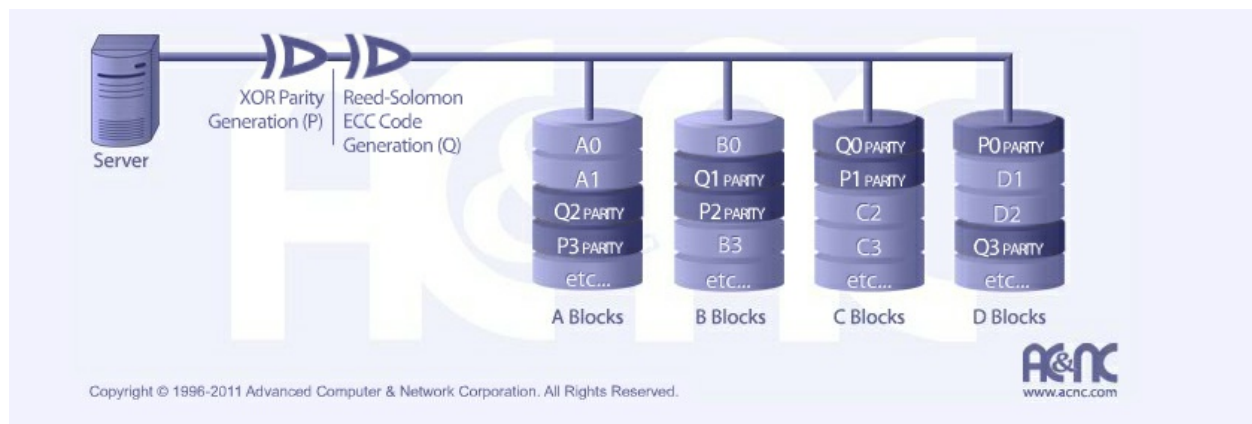


Figure courtesy AC&NC, www.raid.com.

Figure 18.9. RAID-6

RAID-10 (see [Figure 18.10](#)) is a striped array where each segment is a RAID-1 array. Therefore, it provides the same fault tolerance as RAID-1. A high degree of performance and reliability can be delivered by RAID-10, so it is very suitable for high-performance database processing. However, RAID-10 can be very expensive. RAID-10 can be an ideal solution for database implementations requiring high performance and fault tolerance.

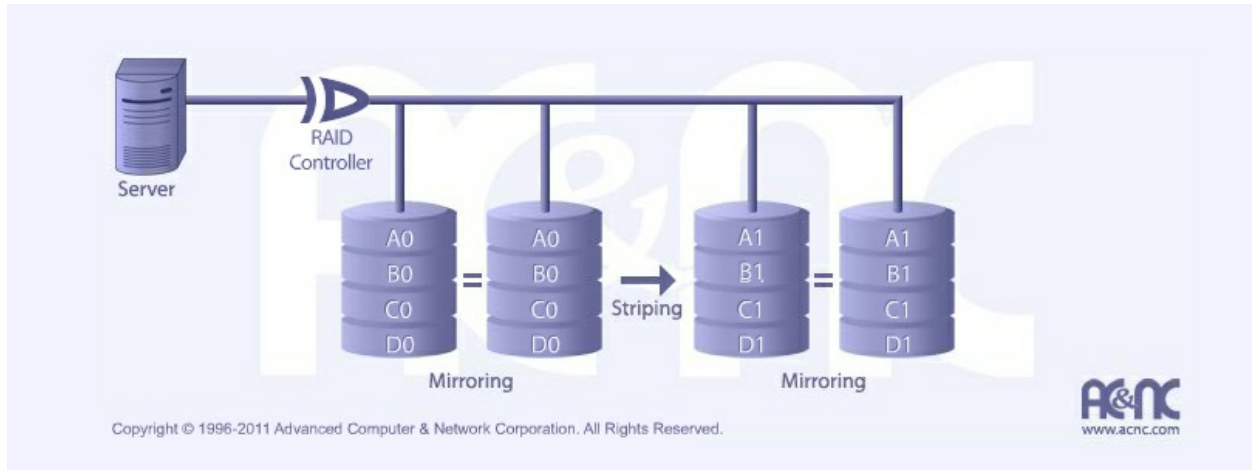


Figure courtesy AC&NC, www.raid.com.

Figure 18.10. RAID-10

RAID-50 (see [Figure 18.11](#)) is a striped array where each segment is a RAID-3 array. RAID-50 is more fault tolerant than RAID-5 but has twice the parity overhead. RAID-50 can achieve high data transfer rates and high I/O rates for small requests, but it is quite expensive to implement.

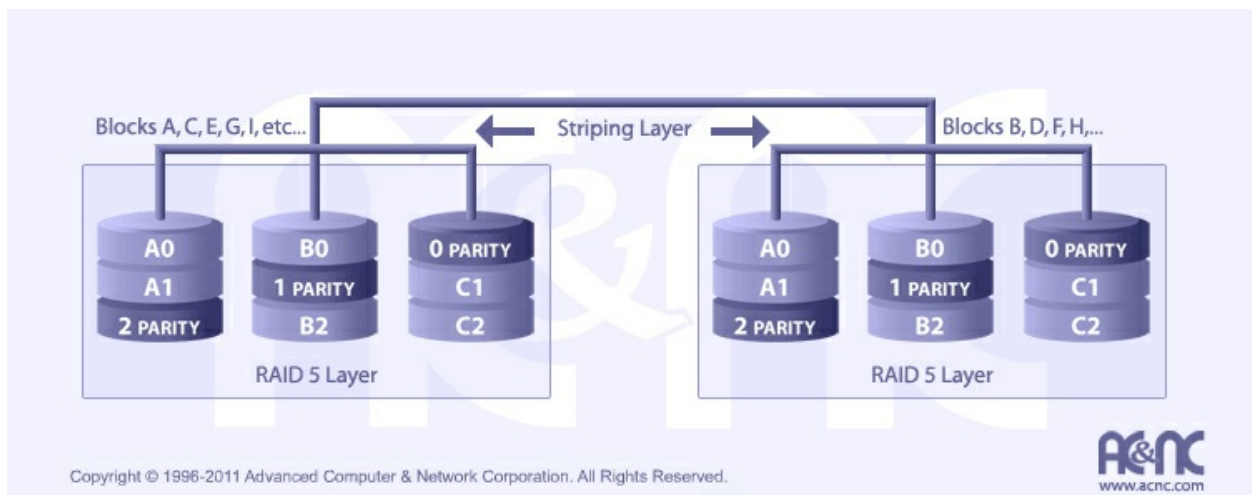


Figure courtesy AC&NC, www.raid.com.

Figure 18.11. RAID-50

RAID-0+1 (see [Figure 18.12](#)) combines the mirroring of RAID-1 with the striping of RAID-0. This couples the high performance of RAID-0 with the reliability of RAID-1.

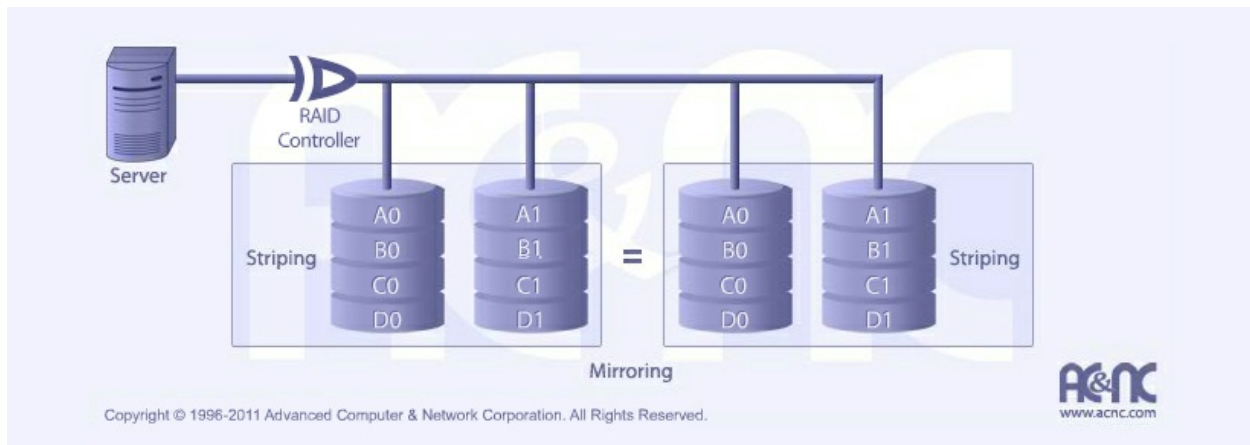


Figure courtesy AC&NC, www.raid.com.

Figure 18.12. RAID-0+1

A number of proprietary variants and levels of RAID have been defined by storage vendors. If you are in the market for RAID storage, be sure you understand exactly what the storage vendor is delivering. For more details, check out the detailed information at <http://en.wikipedia.org/wiki/RAID>.

Proprietary variants and levels of RAID have been defined by storage vendors.

Choosing the appropriate type of RAID storage for database files is an important aspect of the DBA's job. The level of RAID chosen will directly impact the performance and reliability of your database solution. When evaluating RAID storage solutions for your database files, it is essential that you factor the cost of the solution into the selection criteria. [Table 18.2](#) compares the fault tolerance, performance, and cost of each level of RAID.

Table 18-2. RAID Levels and Performance

Level	Fault Tolerance	Read Performance	Write Performance	Cost
No RAID	No	Normal	Normal	Inexpensive
Level 0	No	Fast	Fast	Expensive
Level 1	Yes	Normal	Normal	Moderate
Level 2	Yes	Normal	Normal	Moderate
Level 3	Yes	Normal	Normal	Moderate
Level 4	Yes	Normal	Slow	Moderate
Level 5	Yes	Fast	Slow	Expensive
Level 6	Yes	Fast	Slow	Expensive
Level 10	Yes	Fast	Normal	Expensive
Level 50	Yes	Normal	Normal	Expensive
Level 0+1	Yes	Fast	Fast	Very Expensive

Keep in mind that every DBMS uses files for data, indexes, and transaction logs, each of which serves a different purpose and sustains a different level of read and write activity. Different types of storage can be budgeted and procured for the different types of files and the different types of access required for each type. For example, the write activity against a log file tends to be sequential, so it makes sense to isolate log files on a disk system that maximizes write performance. Because log file activity is primarily write activity, you should favor fast-write RAID levels such as RAID-0+1 or RAID-0. RAID-0+1 is optimal, but quite costly, so many organizations decide to live with the reduced write performance and fault tolerance of RAID-1.

When evaluating RAID—or any type of disk system—you must balance your available budget against the reliability and performance requirements of your database systems and applications. The optimal disk configuration for your database files might require multiple RAID arrays at different levels. For evaluation purposes, though, consider the following advice:

- Favor fault-tolerant RAID levels for database files. Database files not on fault-tolerant disks are subject to downtime and lost data.
- Choose the appropriate disk system for the type of activity each database object will experience. For example, you might want to implement two separate RAID systems—one at RAID-5 for data that is heavily read focused, such as analysis and reporting, and another at

RAID-1 (or RAID-0+1) for transaction data that is frequently written and updated.

- For high-performance, mission-critical implementations with sufficient budget consider RAID-10 for its performance and fault tolerance.
- If you have the budget at your disposal, consider RAID-0+1 because it has fast read, fast write, and fault tolerance.

JBOD

JBOD, which stands for *just a bunch of disks*, is a term used to differentiate traditional disk technologies from newer storage technology. Typically, disks are directly connected to a server. Once the DBMS is installed on those disk drives, database files can be created on those disk drives using the facilities of the DBMS.

JBOD stands for *just a bunch of disks*.

As the database environment becomes more complex, though, the storage needs become more complex. As the database environment grows, DBAs may need to consider other storage techniques such as SANs and NAS.

Storage Area Networks

A *storage area network*, or SAN, generally refers to an interconnected network of storage devices. However, no industry-wide standard definition of SAN exists, and it means different things to different folks. To some, a SAN is anything that includes a fiber channel switch. Others define a SAN to be anything with two or more host systems using fiber channel technology. However you define them, SANs offer high speed, coupled with high availability.

No industry-wide standard definition of SAN exists.

Fiber channel is a serial interface that can deliver a transfer rate of up to 105MB/second. Fiber channel is a competitive technology to *SCSI*, which stands for *small computer system interface* and is pronounced “scuzzy.” SCSI is a parallel interface that can deliver transfer rates from 40MB/second to a maximum of 160MB/second for Ultra-SCSI. With SCSI, though, repairing or adding storage devices causes significant downtime. Fiber channel offers more fault tolerance by enabling devices to be removed and added without

disabling connections.

SAN affords the following benefits:

- Shared storage among multiple hosts
- High I/O performance
- Server and storage consolidation

Network-Attached Storage

Network-attached storage, or NAS, refers to storage that can be accessed directly from the network. With NAS, hosts or client systems can read and write data over a network interface (such as Ethernet, FDDI, or others) using protocols including NFS and CIFS. NAS can be implemented using file servers, host systems running NAS software, or thin servers called appliances.

NAS provides the following benefits:

- Shared storage among multiple hosts
- Simpler management due to reducing duplicate storage
- Application-based storage access at file level

Traditional systems use either internal or external storage attached directly to a system, or perhaps use shared storage, such as a large enterprise array. Performance and distance limitations exist with older interfaces (like SCSI and ESCON). The most commonly used NAS software is NFS (Network File System) and CIFS on Windows and Linux systems.

SAN versus NAS

A SAN is best used as a storage backbone providing basic storage services to host systems and NAS servers. SANs are well suited for sharing storage and building the infrastructure for server and storage consolidation. Applications requiring high performance or large capacities are good candidates for SAN technology. Because it offers high performance and low latency, SAN is ideal for database applications.

NAS is better suited for solving multimedia storage problems, data-sharing issues, and sharing of storage for smaller systems. NAS does not efficiently handle the block-based storage used by database systems.

Do not make the mistake of forcing all storage to be of one type, either SAN or NAS. Match the storage requirements to the access and modification

needs of each database and application. You can use a combination of SAN and NAS to solve your data storage and business requirements.

Don't make the mistake of forcing all storage to be of one type.

Tiered Storage

The final option available for storing database files is tiered storage. With tiered storage, different categories of data are assigned to different types of storage media in order to reduce total storage cost. This approach is important in organizations that manage significant amounts of data that continues to grow in volume. The more data that must be managed by the active system, the more expensive it becomes to manage. Tiered storage can offer some financial relief.

To take advantage of tiered storage, data first must be categorized based on criteria related to the data and its use. Example of categories include

- Usage patterns
- Importance to the organization
- Level of protection required
- Recovery needs (local and disaster)
- Performance requirements

Once categorized, the data can be deployed on the most cost-effective storage medium for its type. Such an approach is sometimes referred to as *information life cycle management*, or ILM. But there are many issues that need to be resolved to make such an approach effective:

- A useful categorization technique needs to be defined.
- A method of measuring data such that it can be categorized is required.
- All of the available storage options need to be aligned with the data categories.
- A technique needs to be created for moving the data to the appropriate storage devices.

Since assigning data to particular media may be an ongoing and complex activity, some storage vendors provide software for automatically managing the process based on user-defined policies.

One approach to categorization, popularized by Teradata, is known as multi-temperature data.⁸ This technique deploys four categories: hot, warm, cool, and dormant. The temperature of data is defined as a function of the access rate for queries, updates, and data maintenance. In other words, hotter data is accessed more frequently than warm data, which is accessed more frequently than cool data. Finally, dormant data is rarely updated or queried and is part of a static data model.

The next step is to categorize your storage devices for use with each data temperature. Hot data that is I/O intensive and requires high availability can be placed on storage devices offering high performance, reliability, advanced features, and large capacity. This typically means high-performance disk drives such as double-parity RAID or perhaps solid-state storage. If the data is not stored on SSD, it should be cached so that it is available in memory. In general, hot data should be stored on the fastest storage technology available and optimized for frequent access.

Warm data is less frequently accessed than hot data and often is read more than it is modified. Such data can be placed on less expensive disk devices that offer good performance and reliability but are not top-of-the-line. Storage devices for warm data should be lower cost than for hot data but with better performance and accessibility than for cool or dormant data. SATA- and SCSI-configured storage can work well for warm data.

Cool data is not dormant but is not accessed often. Such data usually still needs to reside on direct-access storage devices. For such data consider using the lowest-cost disk devices available, perhaps NAS or object-based storage.

Finally, dormant data, which has not been accessed for a long time (perhaps years) and whose data model is stable, can be moved to offline storage systems such as intelligent tape or optical disc.

Planning for the Future

Most database implementations are anything but static. Once deployed, databases are queried, updated, loaded, unloaded, and reorganized, and data is deleted and inserted from them on an ongoing basis. As the data composition of a database changes, its storage requirements will change as well.

The DBA must be ever vigilant in planning for future growth.

The DBA must be ever vigilant in planning for future growth. The DBA must keep an eye on the amount of data and the number of users accessing the data. When either expands, database storage may have to be modified.

Capacity Planning

Capacity planning is a process whereby the capacity of the entire system is measured and compared against requirements. The goal of this comparison is to adjust the resources available to the system as appropriate. To successfully embark on a capacity-planning exercise requires an understanding of new corporate initiatives and how they are likely to impact the existing infrastructure—that is, your hardware (CPU, storage, network, and memory) and your software (application and system).

Capacity planning measures and compares system capacity against requirements.

By measuring current capacity, gauging the growth of capacity over time, and factoring in the anticipated capacity requirements of new corporate and IT initiatives, you can determine whether your existing infrastructure can sustain the anticipated workload. If the projected growth outpaces the ability of your computing environment to support it, you will need to evaluate the cost of modifying and possibly scaling up your computing infrastructure.

From a storage perspective, this may involve simply adding more disk devices and assigning them to the DBMS. However, it may involve additional tasks to support additional data and users, such as the following:

- Redesigning applications
- Redesigning databases
- Modifying DBMS parameters
- Reconfiguring hardware components
- Adjusting software interfaces

There are multiple perspectives from which to view storage consumption. Each is valid and provides a different view of storage utilization and growth. A system-wide perspective views the rate at which disk space is being consumed on an aggregate level. At a lower level, viewing storage capacity by server determines which computers are consuming disk space at the fastest rate. You can also monitor storage consumption at the file system level to

determine which file systems are consuming disk space at the fastest rate. Finally, as a DBA, you will want to view just those files that are associated with the DBMS. The bottom line is that you will need to be able to answer all of the following questions:

- When will more storage be required?
- How much additional storage is needed?
- Where is the additional storage needed?
- What needs to be done to align the additional storage with the DBMS?

Summary

Databases are composed of data and files that must be stored on reusable media to be useful. The DBA must understand the different storage options available for database objects and files and use the appropriate device for each file that meets capacity, performance, and budgetary objectives.

Storage is a growing field.

Storage is a growing field, with new technologies and products becoming available on a regular basis. DBAs must keep abreast of this technology so that they can meet the storage needs of their databases—and thereby meet the information technology and data-processing needs of their organization.

Review

1. Why are most database files stored on disk devices instead of other storage media?
2. What is the difference between SAN and NAS?
3. Name five goals to consider when building a storage system.
4. Your organization has decided to purchase RAID arrays to support your DBMS. Which levels of RAID would you advise them to use and why?
5. What storage management actions should be taken to assure the viability and speed of the transaction log file(s)?
6. Describe the pros and cons of using RAID-4 for database files.
7. When calculating record size for a table, how should you treat the size calculation for variable-length columns?

8. Why is it beneficial to place data files and index files on separate disk devices?
9. What advantages does a SAN have over SCSI devices?
10. Calculating the amount of storage required for a table is as simple as adding up the number of bytes for the columns in a row and multiplying by the total number of rows: true or false?

Suggested Reading

Allspaw, John. *The Art of Capacity Planning*. Sebastopol, CA: O'Reilly (2008). ISBN 978-0-596-51857-8.

EMC Education Services. *Information Storage and Management: Storing, Managing, and Protecting Digital Information*. Indianapolis, IN: Wiley Publishing (2009). ISBN 978-0-470-29421-5.

Lightstone, Sam, et al. *Physical Database Design*. San Francisco, CA: Morgan Kaufmann (2007). ISBN 978-0-12-369389-1.

Poelker, Christopher, and Alex Nikitin. *Storage Area Networks for Dummies*. Indianapolis, IN: Wiley Publishing (2003). ISBN 0-7645-2480-1.

Schur, Stephen G. *The Database Factory*. New York, NY: John Wiley & Sons (1994). ISBN 0-471-55844-3.

Strauss, Melvin J. *Computer Capacity: A Production Control Approach*. New York, NY: Van Nostrand Reinhold (1981). ISBN 0-442-26243-4.

Thornburgh, Ralph H., and Barry J. Shoenborn. *Storage Area Networks*. Upper Saddle River, NJ: Prentice Hall (2000). ISBN 0-13-027959-5.

Toigo, Jon William. *The Holy Grail of Data Storage Management*. Upper Saddle River, NJ: Prentice Hall (2000). ISBN 0-13-013055-9.

———. *The Holy Grail of Network Storage Management*. Upper Saddle River, NJ: Prentice Hall (2004). ISBN 0-13-148968-2.

19. Data Movement and Distribution

Data is not sedentary. Once data has been created, organizations tend to move it around to support many different purposes—different applications, different geographies, different users, different computing environments, and different DBMSs.

Data is not sedentary.

Rarely is a single copy of any piece of data good enough. Data is copied and transformed and cleansed and duplicated and stored many times throughout the organization. Different copies of the same data are used to support transaction processing and analysis; test, quality assurance, and operational systems; day-to-day operations and reporting; data warehouses, data marts, and data mining; and distributed databases. Controlling this vast sea of data falls on the DBA.

The DBA uses many techniques and technologies to facilitate data movement and distribution. This chapter discusses some of the primary tools that are used to move data from place to place and to support data at multiple locations.

Loading and Unloading Data

One of the simplest ways for the DBA to move data from one place to another is to use the LOAD and UNLOAD utilities that come with the DBMS. The LOAD utility is used to populate tables with new data, and the UNLOAD utility is used to read data from a table and put it into a data file.

Again, it is important to note that each DBMS may call the actual utilities by different names, but the functionality is the same or similar from product to product. For example, Microsoft SQL Server and Sybase provide the BCP utility, and Oracle offers SQL*Loader.

The LOAD Utility

A LOAD utility is used to perform bulk inserts of data into database tables. It typically can support

- Adding rows to a table, retaining the current data, or

- Replacing all existing rows with the new data

A LOAD utility is used to perform bulk inserts of data into database tables.

When loading data, the DBA must take into account many factors to ensure success. Is the LOAD utility restartable in the event it fails? Although it takes more time to implement a restartable load process, it is easier to support. For the load process to be restartable, the DBA must ensure that any work files used are allocated and that the LOAD utility can restart from where it left off. If the LOAD utility is not restartable and an error occurs that stops the load process, the DBA must choose one of two options:

1. To delete the data that has already been loaded and start again from the beginning, or
2. To determine what was already loaded and remove those records from the input data file that is being loaded

Of course, this task is more difficult when adding rows to existing data using LOAD, because the first option requires a selective delete.

Another consideration for the DBA is whether to presort data before it is loaded into the database. If the table being loaded is clustered, it may make sense to sort the input data by the clustered columns to allow the LOAD utility to process the input data file in clustered order. Failing to sort data into clustering sequence may cause the load process to uncluster the data, depending on the DBMS and the LOAD utility.

Most database LOAD utilities do not cause triggers to fire, which can cause data integrity problems. If the database and application rely on the trigger to calculate data or enforce constraints, loading data without firing triggers can be devastating to the state of the database. The DBA and application developers will need to develop programs or scripts that mimic the triggered action if data is to be loaded regularly and triggers exist.

Other database constraints can pose problems for the LOAD utility, too. If data is being loaded that does not conform to a unique constraint or a check constraint, what will the LOAD utility do? Once again, this depends on the DBMS and the LOAD utility. Some LOAD utilities are flexible enough to provide options to enforce constraints, or not. In such a scenario, the DBA must decide whether invalid data is to be discarded by the LOAD utility, or

whether data that does not conform to constraints can be loaded—possibly causing data integrity problems. Some LOAD utilities do not even give the DBA an option—they either load everything or discard nonconforming data automatically. All in all, it is better to have the flexibility, even if it does cause more work for the DBA.

The DBA should favor enforcing constraints when the LOAD utility allows it. Otherwise, you will have to correct the data after it has been loaded or live with data integrity violations. It is more efficient to process the data once, as it is loaded, than to process the data twice—once to load it and once to correct it.

Enforce constraints using the LOAD utility.

Before LOAD can be used, the DBA must ensure that all processes and users needing to execute the LOAD utility have the appropriate authority to do so. Some DBMSs provide a specific LOAD authority that can be granted. Most DBMSs permit the owner of the table and those granted administrator-level authority to load data as well.

Describing the Input File

Using the LOAD utility to populate a table requires an input file containing the data. The DBA must define the layout of the input file to the LOAD utility to enable it to translate raw data into the format required for storage in the database. This is typically accomplished by specifying the data type of each column along with the beginning and ending positions within the file where the data exists. Some LOAD utilities also provide specific formats that can be loaded with limited specification such as comma-delimited files or files created using an UNLOAD or EXPORT utility.

The LOAD utility must be capable of handling nulls. Nulls are usually handled with an indicator byte and a specific clause for checking that byte. For example, the LOAD utility could specify a clause such as

[Click here to view code image](#)

```
LOAD . . . STATUS POSITION (20:25) CHAR(6)
NULLIF (26) = 'N' . . .
```

The LOAD utility must be capable of handling nulls.

This statement indicates that the column named STATUS is to be loaded

with data from the input file, starting with the twentieth character and finishing with the twenty-fifth character. However, if the twenty-sixth character is equal to the value N, the STATUS column is to be set to NULL.

Loading tables with nullable columns can degrade the performance of the LOAD utility because extra processing is required to check whether the data should be null or not. If data is to be loaded into a table on a regular basis, try to avoid nullable columns for those tables.

Some LOAD utilities provide additional controls to allow the DBA to modify the functionality of the load process. For example, clauses to control the loading of default values or to bypass certain records in the input file are common. Be sure to understand all of the parameters of the LOAD utility being used so that you can control the way in which data is loaded into your tables. Sometimes the LOAD utility can be used instead of an application program to insert data—which reduces the programmer’s effort because there is no need to write, test, and debug program logic. Instead, the logic exists in the LOAD utility, and the processing can be controlled by the DBA with the parameters and clauses of the utility.

Try to avoid converting data from one data type to another during the load process. Although most LOAD utilities can automatically convert similar data types, additional CPU effort is required to perform such conversions. For example, the LOAD utility may be able to convert an integer value in the input file to a decimal value in the database, but it would be better if the LOAD data were decimal to begin with.

Special care must be taken when floating-point data must be loaded into tables. Usually the LOAD utility will require specific information about the format of the data.

Take special care when loading floating-point data into tables.

Efficient Loading

It is usually a good idea to create all required indexes before loading data into a table. The LOAD utility is usually more efficient at populating the indexes during the load process than creating new indexes for a fully populated table. Of course, the DBA should verify this to be the case for the DBMS and version in use.¹

Create all required indexes before loading data into a table.

If the LOAD utility is capable of performing tasks in parallel, the DBA should take advantage of this when large amounts of data are being loaded. The LOAD utility might be capable of accepting multiple input files for concurrent loading into different segments or table partitions, or it might be able to build multiple indexes in parallel rather than building each sequentially. Parallel operations like this may increase the amount of CPU required to load data while reducing the overall elapsed time for the LOAD to run.²

Another performance-related aspect of loading data is the need to plan for concurrent processing if you are loading multiple tables. Judicious scheduling of concurrent workloads can enable some LOAD utilities to operate concurrently. For example, you might be able to perform concurrent loads on tables from different databases or filegroups. The DBA should make every effort to understand the capabilities and limitations of the LOAD utility and to plan concurrent load jobs accordingly.

If the LOAD utility provides an option to turn off logging, consider using it to speed up the load process and minimize overhead. However, the DBA will have to back up the data after the LOAD completes to ensure recoverability. If logging is turned off, the DBMS cannot recover the table without an image copy backup.

It is usually more efficient for the LOAD utility to insert the initial data into a table than it is for a program or multiple SQL INSERT statements to perform the same task. A LOAD utility is optimized for very efficient bulk insertion of data. In addition to the efficiency of LOAD, using a utility is likely to be less error prone than a corresponding application program.

For some DBMSs, it is more efficient to use LOAD instead of a mass delete. A mass delete is performed when a SQL DELETE statement is issued with no WHERE clause. Using the LOAD utility with an empty input file will have the same result—all rows will be deleted from the referenced table. LOAD is likely to be more efficient, especially if it is possible to turn off logging during the load process.

Running Other Utilities during a LOAD

Some database LOAD utilities take advantage of the fact that they have access to all of the data as it is being loaded. A full-function LOAD utility may enable the DBA to create image copy backups or to create database

statistics during the load process. Taking advantage of such capabilities creates efficiencies because the data is read once but processed multiple times (load, image copy, statistics, etc.).

Additionally, you may want to schedule a REORG after loading data. This is an especially good idea if the loaded data was not presorted into clustering sequence and the LOAD utility does not sort for clustering.

You may want to schedule a REORG after loading data.

The UNLOAD Utility

Information in database tables frequently needs to be moved or copied to other locations. For example, you may want to move data to a different database, from a table to a sequential file for external processing, or possibly to another relational database system or platform. Certain database schema changes require database objects to be dropped and recreated—and when the objects are dropped, so is the data. Therefore, you need to unload the data before making database object changes. Perhaps you just want to extract a subset of rows from a table for use as test data. Even to reorganize a database object typically requires the data to be unloaded, optimized, and then reloaded.

The purpose of the UNLOAD utility is to read data from a database and write it to an output data file. Without an UNLOAD utility, database users are forced to use SQL SELECT statements issued by an interactive SQL facility, report writer, or application program in order to unload data. However, these methods are error prone and slow for large quantities of data. Furthermore, requiring a developer to code an application program to create a file is too inflexible and time-consuming for most production database needs. Thus, many DBMSs provide an UNLOAD utility to supply high-speed capabilities and the flexibility to perform most of the bulk data movement tasks performed by DBAs.

The purpose of the UNLOAD utility is to read data from a database and write it to an output data file.

Before UNLOAD can be used, the DBA must ensure that all processes and users needing to execute the UNLOAD utility have the appropriate authority to do so. Typically, this involves granting a level of authorization that

contains the ability to read (SELECT) data from the table. Depending on the DBMS and utility, additional security may be required.

Concurrency

The level of concurrent activity that can be performed against a table that is being unloaded depends on the parameters and options of the UNLOAD utility. Concurrent read activity is permitted because UNLOAD does not change any data. Most UNLOAD programs enable the user to control whether concurrent data modifications can occur while data is being unloaded. Frequently it is desirable to disable concurrent modifications during an UNLOAD to ensure the creation of a file with consistent data.

Additionally, it may be useful or necessary to prohibit concurrent utilities from running while data is being unloaded. For example, loading, reorganizing, or recovering data during an UNLOAD will cause unpredictable results because these utilities change data.

Unloading from Image Copy Backups

Modern UNLOAD utilities are capable of unloading data from an image copy backup. Such a capability is useful because it enhances concurrent data access. Unloading from an image copy backup can be beneficial because the live data is unaffected—meaning no locks are taken on the live data nor is any data read from the actual table on disk. Because the UNLOAD utility reads the data from an image copy backup, the performance and availability of applications running against the live data will be unaffected by the concurrent unload operation.

A modern UNLOAD utility provides the ability to unload from image copy backups.

Of course, the freshness of the data being unloaded may be an issue. If subsequent updates, inserts, and deletes were processed against the table after the image copy backup was taken, those modifications will not be captured in the unloaded data because they were not made to the image copy data set—only to the live data itself.

If you are in the market for an UNLOAD utility, be sure it provides the ability to unload from image copy backups.

Generation of LOAD Parameters

It is common for an UNLOAD utility to provide an option for generating the control statement necessary to use the LOAD utility to reload the data that is being unloaded. Such an option can save time because the DBA will not have to create the LOAD statement by hand. Even if the data is going to be loaded into a different table, it can be useful to generate the LOAD statement during the unload operation. Changing a generated LOAD statement usually is easier than coding a LOAD statement from scratch.

Data Encoding Scheme

An UNLOAD utility should allow the DBA to specify the encoding scheme to use for the unloaded data. Several common formatting options include EBCDIC, ASCII, and Unicode. Many DBMSs support multiple code sets for other encoding schemes. Depending on the DBMS and the UNLOAD utility, you may be able to specify another data encoding scheme to be used for the unloaded data.

Floating-Point Data

As with the LOAD utility, special care must be taken when unloading floating-point data. When floating-point data is being unloaded, special parameters may need to be specified to identify the format in which unloaded floating-point numbers should be stored.

Limiting UNLOAD Data

Sometimes DBAs will need to unload only a portion of the data. That is, the UNLOAD output file will contain only a subset of the rows in the table. Numerous reasons and situations exist where it might make sense to unload only a subset of the total rows, for example, to create test data or to inspect only certain rows.

It might make sense to unload only a subset of the total rows.

Most UNLOAD utilities provide options to specify a limited number of rows to unload. These options typically take three forms: LIMIT, SAMPLE, and WHEN.

A LIMIT parameter is used to limit the number of rows to be unloaded by the UNLOAD utility. Specifying a parameter such as “LIMIT 200” will cause the UNLOAD utility to stop after unloading 200 rows.

A SAMPLE parameter is used to unload a sampling of the data instead of

the entire table. The SAMPLE parameter typically requires a decimal condition that specifies the percentage of rows to be sampled. For example, the following parameter indicates that 25.75 percent of the rows in the table should be unloaded: SAMPLE 25.75.

Finally, a WHEN clause is used to supply SQL predicates to the UNLOAD utility such that only certain data is unloaded. For example, the following condition will cause only rows where the SALARY is greater than \$50,000 to be unloaded: WHEN (SALARY > 50000). The UNLOAD utility may also provide the ability to use the ORDER BY clause to sort the data as it is written to the unloaded file.

Of course, the exact syntax of these three options will differ from DBMS to DBMS and utility to utility, but the general concept remains the same. Each of these options provides the DBA with great flexibility when creating external data files using the UNLOAD utility.

The exact syntax of the LOAD utility differs from DBMS to DBMS.

Unloading from Views

Most UNLOAD utilities permit data to be unloaded from views, not just tables. Being able to unload data from a view comes in handy when it is necessary to access data from multiple tables and place it into a single file. By creating the view as a join of those tables and unloading from the view, you can unload data from multiple tables to a single file.

Maintaining Application Test Beds

LOAD and UNLOAD can be used to maintain a consistent test bed of data for the purposes of testing application programs, as shown in [Figure 19.1](#). When an UNLOAD data file is created that can be loaded prior to each program test, developers can be sure that each program execution runs against the same data—which is critical for tracking down bugs and assuring that the program code is correct. The DBA can set up the proper LOAD and UNLOAD jobs and then turn them over to the application development teams for their use.

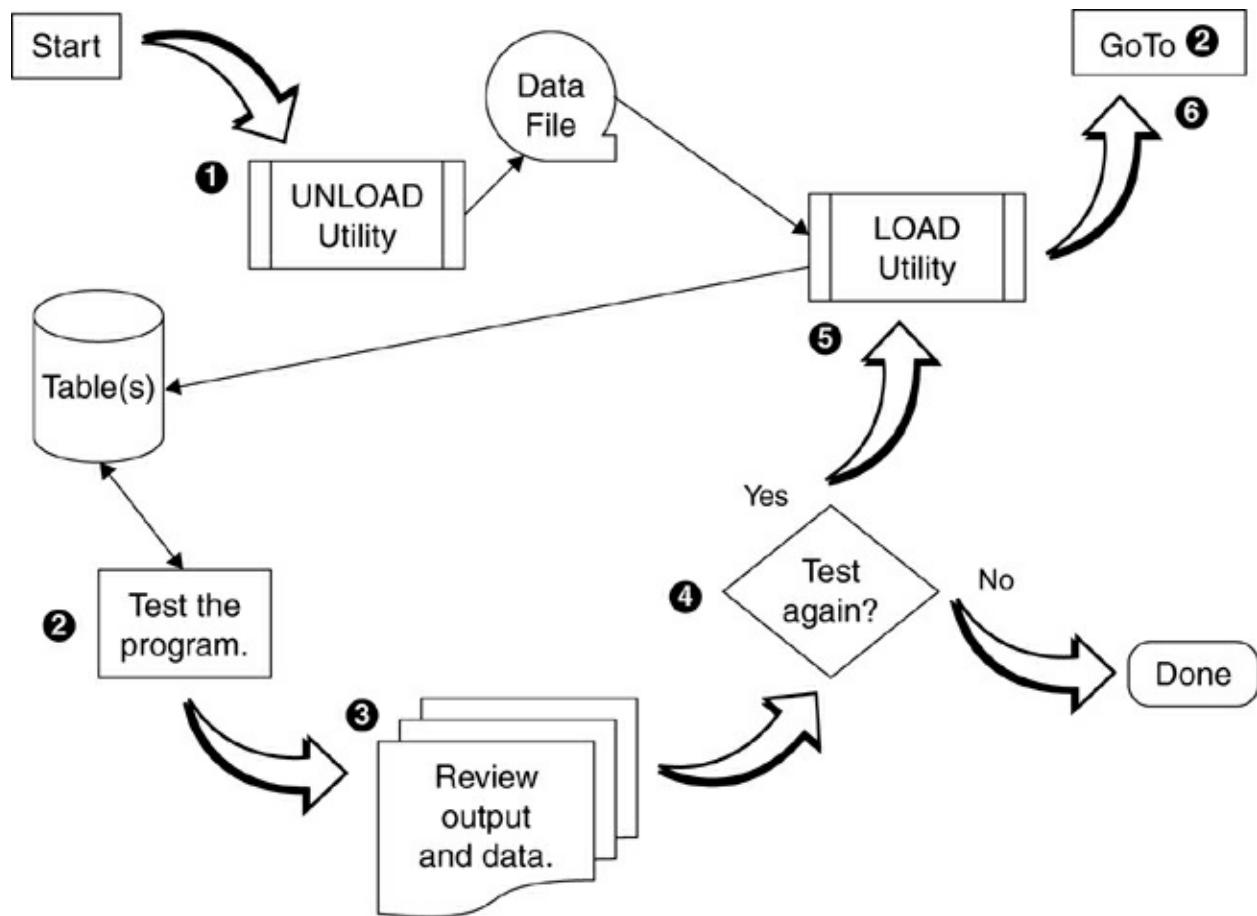


Figure 19.1. Using LOAD and UNLOAD for program testing

LOAD and UNLOAD can be used to maintain a consistent test bed of data.

The UNLOAD utility is quite handy and will be used frequently by most DBAs. There are numerous situations and issues that can be resolved with the UNLOAD utility.

EXPORT and IMPORT

Similar to an UNLOAD utility, an EXPORT utility reads data from a table and places it into an external file. An IMPORT utility reads an external file created by the EXPORT utility and inserts the data into a table.

IMPORT and EXPORT facilities typically work with more than just the data, though. Sometimes an EXPORT data file contains the schema for the table along with the data. In such cases, the IMPORT utility can create the table and import the data using just the EXPORT data file. Sometimes the

EXPORT file contains more than just a single table. Some EXPORT facilities enable the DBA to specify a single table, and then follow the relationships for that table to extract all of the related files and data.

Some IMPORT/EXPORT facilities provide UNLOAD-like features to sample, subset, and limit the data that is exported (and imported). The difference, though, is the ability to perform such functions across multiple tables and maintain referentially intact data.

Not every DBMS offers IMPORT and EXPORT utilities. Some third-party vendors provide import and export products.

Not every DBMS offers IMPORT and EXPORT utilities.

Bulk Data Movement

The combination of UNLOAD and LOAD is the most common method used by DBAs to move large amounts of data from place to place. However, there are other methods for moving large quantities of data. Let's review a few.

ETL Software

ETL is a type of software that performs data movement. ETL stands for *extract, transform, and load*. ETL software is primarily used to populate data warehouses and data marts from other databases and data sources.

ETL stands for extract, transform, and load.

Although most DBMS vendors offer ETL software, it usually is not included with the base DBMS license. Microsoft SQL Server, on the other hand, includes SQL Server Integration Services for ETL with the standard SQL Server license.

By using ETL software, the DBA can automate the extraction of data from disparate, heterogeneous sources. For example, data may need to be extracted from legacy IMS databases and VSAM files from the mainframe; relational databases such as Oracle, SQL Server, and DB2 on various platforms; spreadsheets stored on the LAN; as well as external data feeds. The ETL software can be set up to recognize and retrieve the data from these many different sources.

Once retrieved, the data may need to be transformed in some fashion

before it is stored in the target database. ETL software makes it easy to automate these changes. For example, you may wish to transform coded data to recognizable values, perhaps changing 1, 2, and 3 to “Married,” “Single,” and “Divorced.” In addition, ETL software is capable of changing data types, and aggregating and summarizing numerical data.

After the data is transformed to your specifications, it is loaded into the target database. ETL software is much more flexible and useful for complex data movement needs than simple LOAD and UNLOAD utilities.

Replication and Propagation

Another method of moving data is through [replication](#) and [propagation](#). When data is replicated, one data store is copied to one or more data stores, either locally or at other locations. Replication can be implemented simply by copying entire tables to multiple locations. Alternatively, replicated data can be a subset of the rows and/or columns. Replication can be set up to automatically refresh the copied data on a regular basis.

Propagation, on the other hand, is the migration of only changed data. Propagation can be implemented by scanning the transaction log and applying the results of data modification statements to another data store. Initial population of a data warehouse can be achieved by replication, and subsequent population of changes by either replication (if the data is very dynamic) or propagation.

The difference between replication and propagation is highlighted in [Figure 19.2](#). DBAs can use both of these techniques to move data between different databases and data warehouses. More information on managing and administering databases in a data warehouse environment is provided in [Chapter 20, “Data Warehouse Administration.”](#)

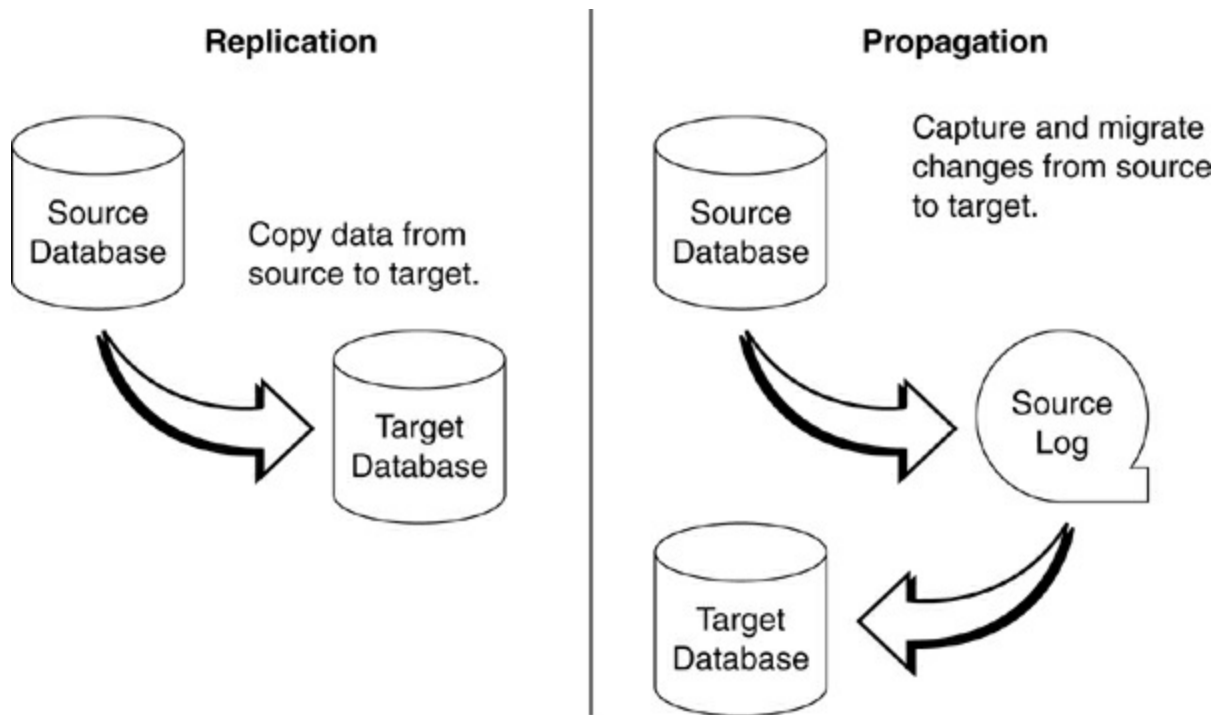


Figure 19.2. Replication and propagation

Messaging Software

Messaging software, also known as *message queuing software* or *application integration*, is another popular form of data movement. When using a message queue, data is placed onto the queue by one application or process; the data is read from the queue by another application or process.

Messaging software works by providing APIs to read and write formatted messages to and from a queue. An application can read or write messages to and from the queue from any platform supported by the software.

Messaging software provides many benefits for moving data, but primarily the benefit is simple, heterogeneous, any-to-any connectivity among disparate platforms from the desktop to mainframes. Messaging allows businesses to easily integrate disparate islands of data in a time-independent manner. As they both use the message queue, two applications can communicate with each other even if they are not running at the same time.

Messaging allows businesses to easily integrate disparate islands of data.

Popular messaging software includes IBM's MQSeries, Microsoft's

Service Broker, and Oracle’s Advanced Queuing feature.

Other Methods

Of course, many other methods exist for moving data—from the simple, such as using a table editing tool to highlight and copy data, to the complex, such as writing programs to read the database and write to external files or directly to other databases.

Some DBMSs provide additional built-in methods for copying and moving data. For one example, see the sidebar “[Oracle Transportable Tablespaces](#).”

Oracle Transportable Tablespaces

Oracle provides the capability to transport an entire tablespace, including its definition and contents, from one database to another. In some cases, transporting a tablespace can be much easier and more efficient than unloading data from one database and loading it into another.

This option requires the same exact database object definitions to exist in each database. Furthermore, the data contents will be the same after you transport the tablespace (after which, of course, you can edit the data).

Transportable tablespaces are quite useful when data must be moved from database to database in an Oracle system. It is usually easier to copy files and run an IMP command on the destination database than it is to use SQL*Loader or a series of export files to accomplish the same thing. Furthermore, neither the source nor the target database instance needs to be shut down in order to transport tablespaces.

Distributed Databases

Sometimes simply moving data from one location to another is not sufficient. Instead, the data needs to be stored at, and accessible from, various locations throughout an organization. In this situation, a distributed database is required.

A *distributed database* permits data to reside at different physical locations in different databases, perhaps using different DBMS software on different operating systems. Consider, for example, an organization with retail outlets that implements a distributed database system. Each retail outlet would have

a database, and the headquarters would house a central database. With networking technology and the distributed capabilities of the DBMS, data could be accessed and modified at any location from any location. Furthermore, you could specify which locations have update, or even read, access to certain databases.

A distributed database permits data to reside at different physical locations in different databases.

In summary, the location of data in a networked environment does require an analysis of many design considerations for fast access time, high integrity, and easy maintainability. The basic premise of distributed database technology is to provide uniform access to data that is logically related but physically distributed in multiple locations. There are multiple types and variations in the way that a distributed database can be implemented.

The degree of database distribution that can be achieved is defined by the technology and capabilities of the DBMS in use. The following characteristics of a distributed environment differ by DBMS:

- *Autonomy* represents the degree to which sites within the distributed database implementation can operate independently. An intermediate level of autonomy exists when each site can voluntarily share its data. In contrast, tight coupling represents a high degree of autonomy because each site in the distributed environment has complete knowledge of the state of the system and can control the processing of user requests that span data at multiple sites.
- *Isolation* defines whether each site is aware of the other sites in the distributed environment or whether each site is a stand-alone server, unaware of the other sites.
- *Transparency* refers to the shielding of the data location from users and applications.

The characteristics of a distributed environment differ by DBMS.

Each DBMS delivers functionality to implement distributed databases with differing degrees of autonomy, isolation, and transparency. To set up an efficient and effective distributed database environment, the DBA must

thoroughly understand the capabilities of the DBMS. Furthermore, the DBA must offer guidance to the application developers. In a system with distribution transparency, the developers need not know the location of data; more typically, distributed programming requires the programmer to explicitly code a connection to a database server location before accessing data from that location.

A distributed database can be set up as a single database or as multiple databases. With the single database setup, a single DBMS process controls all data access. A multidatabase setup, on the other hand, provides multiple independent DBMS processes, each controlling access to its own local database.

A multidatabase setup can be federated or unfederated. A *federated* multidatabase setup spreads control among the distributed databases. The DBA can control each local database component of the federation to determine what data is stored there and who can access it. An *unfederated* multidatabase scheme centralizes the control of these issues. Typically, unfederated multidatabase setups are homogeneous, whereas federated multidatabase setups are heterogeneous.

Setting Up a Distributed Environment

A distributed DBMS allows data to be physically located on more than one DBMS and at more than one location. Truly distributed data can reside on database servers as well as on client nodes. The more distributed the data, and the closer the data is to the requester, the better the performance is likely to be. However, as data is distributed to more locations, it becomes much more difficult to manage and administer. From a maintenance perspective, the administration required for a large number of desktop DBMSs is quite difficult. A real-life distributed database implementation will often consist of data residing at multiple server locations, and sometimes on client nodes, too.

Distributed data can reside on database servers as well as on client nodes.

From the perspective of the DBA, some of the most difficult aspects of implementing a distributed database are designing it and setting it up. The DBA's first task is to understand the distribution capabilities of the DBMS. Can it participate in federated or unfederated systems? What additional

software is required to support distributed databases? Does the DBMS support two-phase COMMIT? What types of networking protocols are required?

The DBA must understand the distribution capabilities of the DBMS.

The DBA will need to work with network administrators to assure that the company's networks are properly configured for supporting distributed databases. The DBMS will require configuration parameters to be set that enable distribution and that outline the network locations for the remote databases.

Up to this point, we have discussed the physical implementation details. But what about the database design aspects of distributed databases? What data goes where—and why? Suppose you are setting up a distributed database such as that depicted in [Figure 19.3](#).




Figure 19.3. A sample distributed database implementation

In our sample distributed database, the data resides in Seattle, Washington, and in Sydney, Australia. How does the DBA decide which data should be stored in Seattle and which data should be stored in Sydney? How does the DBA provide the proper authorization and configuration such that users can access the required data regardless of its location?

Frankly, there are no hard-and-fast rules that answer these questions. There

are, though, many guidelines that can be followed. First, data should reside at the server where it is most likely to be accessed. The DBA should analyze the intended usage of data to optimize application performance by reducing transmission costs. For example, data used most frequently in Seattle should reside on the Seattle database server, and data used more often in Sydney should reside at the Sydney site.

Data should reside at the server where it is most likely to be accessed.

However, data placement decisions are usually not that simple. In many distributed systems, data may reside in two or three locations but be accessible from many other locations. What if a certain subset of data is used most frequently in Singapore? Once again, store that data on the closest database server—in our example, Sydney is closer to Singapore than to Seattle.

An additional concern arises when data usage is spread almost evenly between the two sites. What if data is required 50 percent of the time by Sydney processes and 50 percent of the time by Seattle processes? In such cases the DBA must choose one location over the other or choose to keep replicas at each site. If replication is not an option, place the data at the location that is most easily managed and monitor the performance of all distributed accesses very closely.

As a DBA, you will need to provide usage guidelines for distributed databases to the application development teams. Such guidelines should cover where data is located, the performance impact of accessing distributed data, and how to optimize distributed relational access. Some of these guidelines depend on the DBMS implementation, whereas others, such as the following, are DBMS independent:

- Minimize network traffic by retrieving only the data that is needed; that is, request no columns or rows in addition to what are absolutely required by the application process.
- Use methods that encourage sending data over the network in blocks instead of one row at a time.
- Access local data instead of remote data whenever it is possible to do so.

Data Distribution Standards

There are two common standards supported by the major DBMSs for distributing data: DRDA and RDA. Both DRDA and RDA are similar in their goals.

Distributed Relational Database Architecture, or DRDA, is IBM's protocol for accessing distributed data regardless of its physical location. DRDA provides methods of coordinating communication among distributed locations. This allows applications to access multiple remote tables at various locations and have them appear to the end user as if they were a logical whole. When a DBMS is DRDA compliant, it means that the DBMS follows the DRDA specifications for distributed databases.

DRDA is IBM's protocol for accessing distributed data regardless of its physical location.

A distinction should be made, however, between the architecture and the implementation. DRDA describes the architecture for distributed data and nothing more. It defines the rules for accessing the distributed data, but it does not provide the actual APIs to perform the access. DRDA is not an actual program but a set of standards.

Remote Database Access, or RDA, is a competing set of protocols developed by the ISO and ANSI standards committees. RDA was built to work with a subset of SQL that is available to most DBMSs.

RDA is a set of protocols developed by the ISO and ANSI standards committees.

RDA is used to establish a remote connection between a client and a server. The client interfaces to a process that controls data transfers to and from the database. The goal of RDA is to enable applications to interconnect with heterogeneous databases and environments.

As an alternative to DRDA or RDA, a gateway product can be used to access distributed data. Gateways are composed of at least two components—one for each distributed location—that communicate with each other to permit data to be accessed.

Accessing Distributed Data

The DBA must understand the type of access supported by the DBMS and help the application developers match the DBMS capabilities to application requirements. Many different types of access are possible, but not all are supported by each DBMS, and further limitations exist with heterogeneous implementations.

Not all methods of access are supported by each DBMS.

The simplest type of distributed database access is a *remote request*. This consists of a single request to a single location within a single unit of work. A remote request allows developers to operate within one DBMS while referring to a different DBMS. This approach is the easiest—but least flexible—method of coding distributed access.

A more complex type of distributed access is called *remote unit of work*. It takes place when a single application program accesses data from multiple locations, but not within the same unit of work. The programmer must know where data is located and build units of work that access data by location. In this scenario:

- Each request must be for a single location.
- Each unit of work can contain multiple SQL statements.
- Each unit of work must access data from a single location only.

Thus, multiple SQL requests per unit of work are permitted, but only to one DBMS per SQL request.

The next step is to remove the unit of work limitation so that each unit of work can access data from multiple locations. This type of access is referred to as *distributed unit of work*. In this case, more than one DBMS can be accessed per unit of work. Multiple SQL statements can read and/or modify data at multiple database servers within a single unit of work.

The final and most robust form of distributed access is the *distributed request*, where a single SQL statement accesses data from more than one location at a time. Thus, a SQL statement can access multiple database servers, and multiple SQL requests, both distributed and nondistributed, can be contained within a single unit of work. This is the most flexible level of distributed database access.

Frequently an additional product is required to support a distributed request, such as IBM's DataJoiner. [Table 19.1](#) outlines the different types of

distributed data access, using the names given to them by the DRDA protocols.

Table 19.1. Levels of Distributed Database Support

Distributed Access	SQL Stmts per UOW	DBMSs per UOW	DBMSs per SQL Stmt
Remote request	1	1	1
Remote unit of work	>1	1	1
Distributed unit of work	>1	>1	1
Distributed request	>1	>1	>1

Two-Phase COMMIT

In order to modify distributed data across two different platforms, the DBMS must be able to ensure that the modifications are treated as a single operation. As with any unit of work, all of the SQL statements within the COMMIT scope either succeed or fail. Therefore, for each distributed COMMIT, either the results of all operations are applied to each database, or none of the results are applied—regardless of the database and its location. This requires a two-phase COMMIT, where one phase is the preparation and the other is the COMMIT.

The DBMS must be able to ensure that the modifications are treated as a single operation.

Distributed two-phase COMMIT enables application programs to update data in multiple database servers within a single unit of work. The two-phase COMMIT process coordinates the COMMITs across the multiple platforms. The two-phase COMMIT provides a consistent outcome, guaranteeing the integrity of the data across platforms, regardless of communication or system failures. One of the database servers acts as the coordinator of the two-phase COMMIT, while the other database servers are participants.

During the preparation phase, each participant prepares for a COMMIT. Each participant informs the coordinator when log records are successfully written and indicates that it is ready to COMMIT changes. When all participants are ready to COMMIT, the second phase, the actual COMMIT, begins. This phase is implemented as a series of communications between the

coordinator and its subordinate participants. During the COMMIT phase, success is presumed, even in the case of system failure. Because all participants have elected to continue the COMMIT, success can be presumed with no danger of data integrity violations. However, if any participant fails to COMMIT, the coordinator will need to roll back the changes for all participants.

Whenever applications perform multisite updates within a single unit of work, a two-phase COMMIT must be used to ensure data integrity.

Distributed Performance Problems

Performance can be particularly troubling in a distributed database environment. As with any system that relies on network activity, the biggest threat to performance is network traffic. The more data that must be sent along the network, the greater the potential for performance problems.

The biggest threat to performance is network traffic.

But let's back up for second. Recall the definition of performance given in [Chapter 9](#): "Database performance is the optimization of resource use to increase throughput and minimize contention, enabling the largest possible workload to be processed." Performance in a distributed environment also is defined by throughput and response time. However, the requester and the server emphasize different aspects of this definition.

Performance in a distributed environment is defined by throughput and response time.

The server views performance primarily in terms of throughput. Remember that throughput is the amount of work that can be done in a unit of time. The server must serve multiple requesters. The more requesters that can be serviced, the better the server is performing. The requester, however, views performance more in terms of response time. Response time is the amount of time required to accomplish a predefined set of work. Response time is a more useful indicator to end users because they are the ones who are waiting for a result—the longer the response time, the longer it takes to complete the user's job.

To analyze the throughput of a distributed database request, the DBA must inspect the entire throughput chain required to satisfy that request. Failure to

analyze the performance of any one component can result in performance degradation. The throughput chain for a request includes every piece of hardware and software and every configuration that must be traversed to deliver service to that end user. Typical components of a throughput chain include

- The computer hardware, local operating system, networking software, and local databases for the requester
- The network hardware, wiring, gateways, routers, and hubs
- Any middleware or transaction processing system used by the requester or the server
- The computer hardware, local operating system, networking software, and databases on the server
- Disk storage and storage management software

Each link in the chain is required to complete a given transaction. The best throughput that any given configuration can achieve is always confined by the slowest component of the chain. To tune distributed performance, the DBA should expend more effort on the weaker links in the throughput chain.

Summary

Because data is required at various points throughout an organization, the DBA must understand the various ways of moving data around. This can involve copying data from place to place—either complete copies of the data or a subset of the data. Utilities such as LOAD and UNLOAD, IMPORT and EXPORT, are ideal for accomplishing many of these tasks.

Sometimes the database environment must be designed such that logically related data is distributed to different physical locations. This is a distributed database. The DBA must be capable of designing distributed databases given knowledge of the capabilities of the DBMS, the needs of the organization, and the locations that must be supported.

Data is likely to reside in global locations within your organization. As a DBA, it is your duty to ensure its integrity, efficiency, accessibility, viability, and recoverability.

Data is likely to reside in global locations within your organization.

Review

1. Why is a two-phase COMMIT necessary when data is modified at two locations within a single unit of work?
2. Describe a technique for unloading data from multiple tables into a single UNLOAD file.
3. What techniques can be used to limit the amount of data that is unloaded?
4. What is the biggest threat to efficient performance for a distributed database system?
5. Describe the two-phase COMMIT process.
6. How can LOAD and UNLOAD be used to maintain test beds of data for applications?
7. How are nulls handled when loading data into a table using the LOAD utility?
8. Name and define the three capabilities used to describe distributed technology and databases.
9. What is the difference between replication and propagation?
10. What issues and constructs can cause data integrity problems after a LOAD?

Bonus Question

Read the following paragraphs and describe a distributed database implementation that would make sense for the data discussed. Keep in mind the concepts of data placement for distribution efficiency as you outline where each type of data should be kept.

Acme Corporation is headquartered in Phoenix, AZ. It has distribution centers located in Los Angeles, New York City, London, and Tokyo. Acme wishes to build a distributed database for tracking goods from its distribution centers to its retail outlets.

All edible goods are distributed out of the Los Angeles distribution center, except for meat, fish, and poultry, which are distributed from every location. Electronics are distributed out of New York City only. European goods are distributed out of the London center, and Asia-

Pacific goods are shipped out of the Tokyo center. All other goods can be distributed from any distribution center.

Suggested Reading

Bell, David, and Jane Grimson. *Distributed Database Systems*. Wokingham, England: Addison-Wesley (1992). ISBN 0-201-54400-8.

Bobak, Angelo R. *Distributed and Multi-Database Systems*. New York, NY: Bantam Intertext (1993). ISBN 0-553-09156-5.

Burleson, Donald K. *Managing Distributed Databases*. New York, NY: John Wiley & Sons (1994). ISBN 0-471-08623-1.

Chorafas, Dimitris N., and Heinrich Steinmann. *Solutions for Networked Databases*. San Diego, CA: Academic Press (1993). ISBN 0-12-174060-9.

Hackathorn, Richard D. *Enterprise Database Connectivity*. New York, NY: John Wiley & Sons (1993). ISBN 0-471-57802-9.

Podcameni, Silvio, Manfred Mittelmeir, and Michele Chilanti. *Distributed Relational Database: Cross-Platform Connectivity and Applications*. 3rd ed. Upper Saddle River, NJ: Prentice Hall (1996). ISBN 0-13-570797-8.

Schur, Stephen G. *The Database Factory*. New York, NY: John Wiley & Sons (1994). ISBN 0-471-55844-3.

20. Data Warehouse Administration

The ongoing quest for more and better business intelligence has caused many organizations to develop data warehouses and data marts to serve their analytical needs. Although, at the most basic level, a data warehouse is simply a database that has been designed and optimized for a specific type of processing, DBAs need to administer them differently from typical databases. This chapter outlines why this is so.

A data warehouse is designed and optimized for a specific type of processing.

What Is a Data Warehouse?

So, just what is a data warehouse, and why should you treat it differently from a DBA perspective? William H. Inmon, recognized as the father of the data warehouse, defines [*data warehouse*](#) in his book *Building the Data Warehouse* (1996) as follows:

Data Warehouse:

A subject-oriented, integrated, nonvolatile, and time-variant collection of data in support of management's decisions.

But what exactly does that definition mean? Let's break it down.

- *Subject-oriented.* Data pertains to a particular subject instead of the many subjects pertinent to the company's ongoing operations.
- *Integrated.* Although the data can originate from a variety of disparate sources, the data within a data warehouse is merged into a coherent whole.
- *Nonvolatile.* Data is stable in a data warehouse. Although data can be added, data cannot be removed. This provides a consistent picture of the business for analysis.
- *Time-variant.* Data stored within a data warehouse is identified with a particular time period.

For our purposes, data warehousing can be defined as the process of

extracting, integrating, transforming, and cleansing data and storing it in a consolidated database. The data warehouse contains atomic data and lightly summarized data. Once the data is consolidated and made available, the data warehouse becomes the only data source that management should access for decision making.

Data warehousing is the process of extracting, integrating, transforming, and cleansing data and storing it in a consolidated database.

What, then, is a data mart? A data mart is basically a departmental data warehouse defined for a single business unit or area. Actually, there is no universally agreed-on difference between a data warehouse and a data mart. A data mart is not necessarily smaller than an enterprise data warehouse. It may be smaller, but size is determined based on business needs. For example, departmental analysis at the business unit level may require more historical information than cross-department, enterprise-wide analysis.

At any rate, both data marts and data warehouses need to be treated differently from traditional database systems from a DBA perspective.

Analytical versus Transaction Processing

The key difference between a traditional database and a database designed as a data warehouse is the type of processing each supports. The traditional database is designed primarily for transaction processing—that is, supporting the transactions that are required for business operations to be conducted. Transactions read and write data on a regular basis. In contrast, data warehouses are designed for analytical processing—that is, supporting the business intelligence and knowledge discovery needs of the organization. Analytical operations are read intensive with very few, if any, write operations required.

Data warehouses are designed for analytical processing.

Terms that are commonly used to describe analytical processing include decision support, DSS, business intelligence (BI), advanced analytics, OLAP, data mining, and information center queries. Each of these activities is read intensive, consisting of queries that sift through the data looking for patterns and trends.

DSS stands for *decision support system*, which is a read-only database designed for analytical processing.

OLAP stands for *online analytical processing*. With *OLAP* technology, end users can navigate through a data warehouse to derive intelligence from data through interactive and iterative queries against the warehoused data. *OLAP* uses a multidimensional view of detail, summary, and aggregate data to access information for further analysis. The key term here is *multidimensional*. A dimension consists of structural attributes viewed as similar by the end user. For example, years, quarters, months, days, and so on make up the time dimension.

OLAP uses a multidimensional view of detail, summary, and aggregate data.

A *dimension* is a modifier of the basic fact that is being analyzed. Examples of *facts* include sales amounts, inventory totals, and expense amounts. Each fact depends on multiple dimensions. For example, sales amounts will differ by geographic region, time, product, and so on. *OLAP* permits multidimensional analysis of facts by applying complex calculations across dimensions. Additionally, *OLAP* provides many other capabilities, including analysis and trending over time, subsetting of data, drill-down through varying levels of detail, “reach-through” to operational data, and data comparison methods.

The bottom line is that *OLAP*, in contrast to *OLTP*, focuses on analytical, rather than transactional, processing. [Table 20.1](#) compares *OLAP* and *OLTP* characteristics.

Table 20.1. OLAP versus OLTP

Characteristics	OLAP	OLTP
Scope of work	Entire database	Single transaction
Amount of data	Individual rows	Groups of rows
Nature of data	Derivative	Primitive
Volatility of data	Low—rarely updated	High—regularly updated
Expected response time	Minutes, perhaps hours	Subsecond
Machine usage	Dynamic and unpredictable	Stable and predictable
Processing priorities	User flexibility	High performance and availability
Type of access	Undefined and dynamic	Predefined and static

Data mining is the process of discovering heretofore unknown information and patterns lurking within an organization's data. Data mining is fast becoming a requirement for the modern, competitive company. Data mining processes apply heuristic algorithms to historical data to automatically uncover patterns in the data. The quantity and quality of available data and the caliber of the pattern discovery algorithms determine the value of the data mining applications. Consequently, organizations that deploy data mining tend to store more data for longer periods.

Both OLAP and data mining are methods for searching out patterns in corporate data. However, OLAP is user driven, while data mining is computer driven. This introduction to OLAP and data mining is necessarily brief—an entire book could be devoted to each of these subjects. The key from the DBA perspective is to plan for administering and managing the data warehouse databases in a manner consistent with their usage.

OLAP and data mining are methods for searching out patterns in corporate data.

Information Center is an antiquated term that was used for DSS processing in the early 1980s. Those readers who recall the proliferation of the Information Center in the 80s will associate it with data warehousing.

Administering the Data Warehouse

The issues of importance when managing a data warehouse differ from those of traditional database administration. Some of the tasks are equivalent, but the implementation and expectations are different.

Data warehouses are accessed mostly by read operations for analytical processing, whereas the traditional database is accessed by a mix of read and write operations for transactional processing. The data warehouse is strategic—it allows users to discover hidden trends and new ways of conducting business. The traditional database is tactical—it allows the business to conduct day-to-day operations.

Be aware of the complexity of data warehouses before undertaking any implementation project. Detailed knowledge of the applications accessing the databases that feed the data warehouse must be available. Be sure to allot development time for learning the complexities of the source systems. Furthermore, the documentation for production systems is often inadequate

or nonexistent, so additional time can be lost just trying to understand what already exists. Furthermore, analyzing the source data to determine what level of data scrubbing is required is a time-consuming task.

All of these issues must be factored into the task of data warehouse administration. Let's examine some of the different tasks required to administer data warehouses efficiently and effectively.

Too Much Focus on Technology?

When you're developing a data warehouse, be sure to include tools, people, and methods in your warehouse blueprint. Too often, the focus is solely on the technology and tools. A data warehouse project requires more than just sound technology. It needs careful planning and implementation (methods) as well as a means to learn from the efforts of others (people) through mentoring, consulting, education, seminars, and user groups.

A data warehouse project requires more than just sound technology.

Data Warehouse Design

When designing the data warehouse, be sure to create a physical design with the special needs of the data warehouse in mind, instead of just mimicking the design of similar production OLTP databases. Common physical design techniques used for OLTP databases, such as a database in third normal form, can in fact inhibit an OLAP database from performing well. This is so because the data access requirements for file and table structures, indexes, and clustering sequence in a data warehouse differ from those of OLTP databases, as do the access requirements and patterns.

A common database design for the data warehouse is the [star schema](#). This name derives from the pattern formed by the data model when it is diagrammed, as in [Figure 20.1](#). In a star schema design, a central fact table stores the primary business activity at the center of the star. This fact table is usually based on a numeric value or a group of numeric values. The fact table is surrounded by the dimensions that affect the activity, and the dimension tables make up the points of the star.

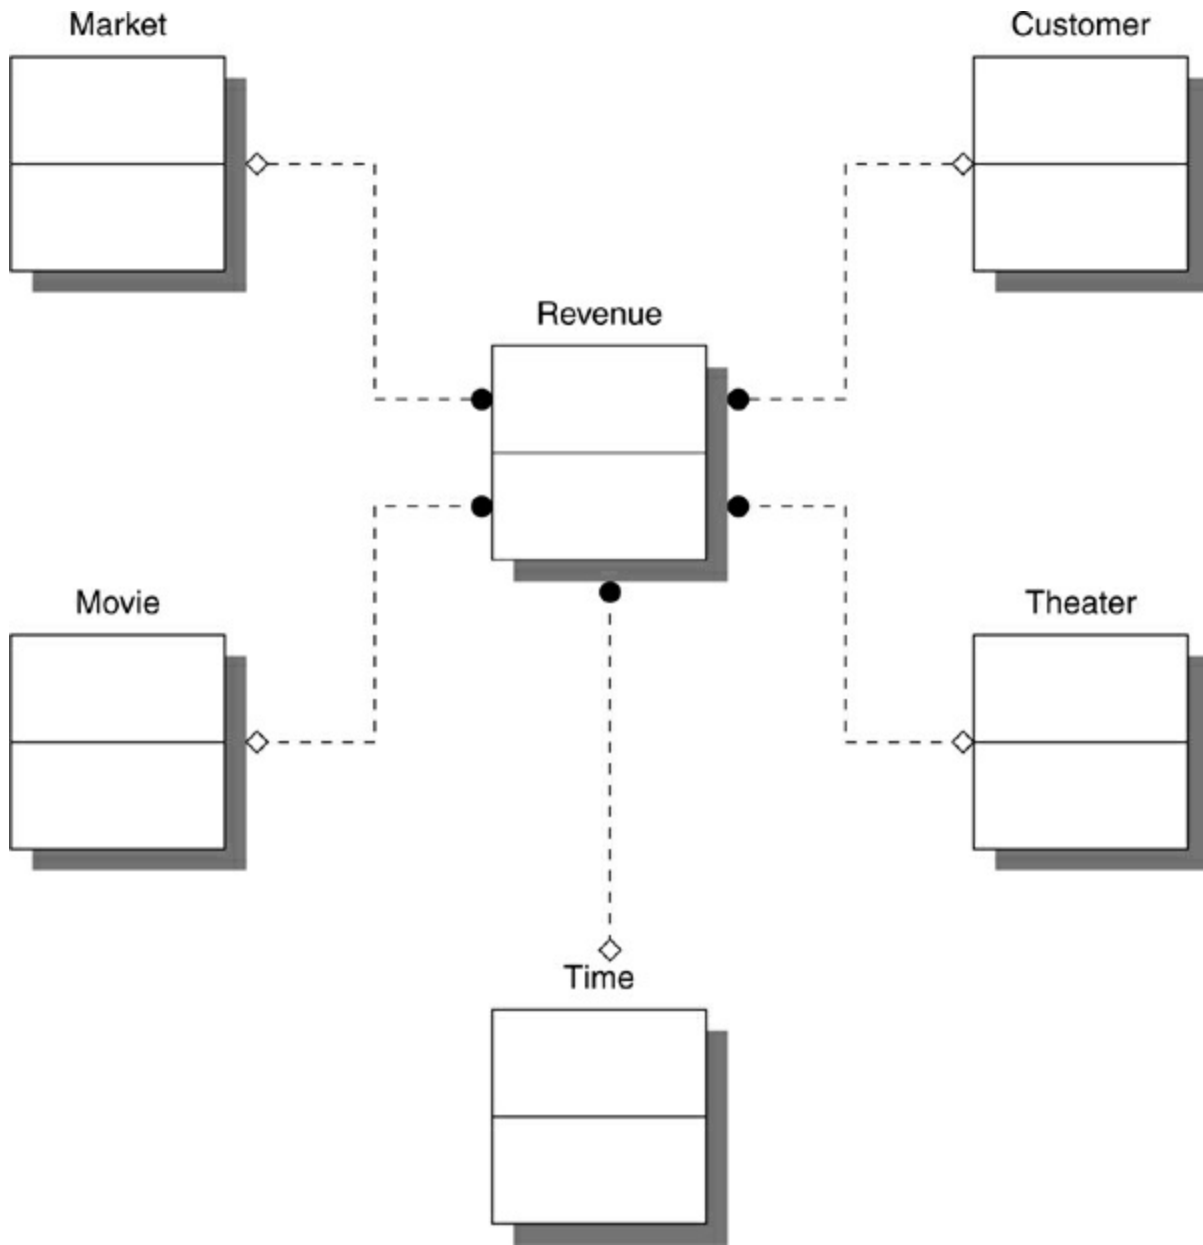


Figure 20.1. The star schema

A common database design for the data warehouse is the [star schema](#).

In this simple example, the Revenue table is the fact table, and the five dimensions of the star are the Movie, Theater, Customer, Market, and Time tables. As an end user, you may want to review revenue by any one, or perhaps a combination, of these dimensions.

The star schema is well suited for the design of data warehouse databases

for the following reasons:

- The design is flexible enough to facilitate quick changes or additions as the data warehouse grows or the application changes.
- The design is easy to understand for both developers and end users.
- The design mimics the way end users typically think about and use their business data.
- The design is easily implemented as a physical database and results in efficient queries because many DBMSs are built to recognize and optimize star joins.

Sometimes additional details are required, and the star schema becomes a snowflake schema. A *snowflake schema* is a star schema in which the dimension tables have additional relationships. When additional tables are required that relate to one or more dimension tables, a snowflake schema is born.

A DBA working on a data warehouse implementation should take care to optimize the DBMS installation for star schema support. Some DBMS products require special installation parameters to enable the optimizer to more efficiently process star schema queries. Be sure to examine the installation parameters carefully and set them properly for data warehousing, otherwise some of your queries may be suboptimally processed.

Of course, star schema and snowflake schema designs are not the only aspects of data warehouse design of concern to the DBA. To aid in the performance of certain regularly occurring queries, it is common for some portions of the data warehouse to be denormalized. Because a data warehouse is a read-only database, the DBA should consider the many denormalization options discussed in [Chapter 4](#). [Table 20.2](#) summarizes the various denormalization options to consider.

Table 20.2. Types of Denormalization

Denormalization	Description
Prejoined tables	Used when the cost of joining is prohibitive
Report tables	Used when specialized critical reports are needed
Mirror tables	Used when tables are accessed concurrently by multiple types of environments
Split tables	Used when distinct groups use different parts of a table
Combined tables	Used to consolidate one-to-one or one-to-many relationships into a single table
Redundant data	Used to reduce the number of table joins required
Repeating groups	Used to reduce I/O and (possibly) storage
Derivable data	Used to eliminate calculations and algorithms
Speed tables	Used to make the processing of hierarchies more efficient
Physical denormalization	Used to optimize for specialized physical DBMS characteristics

As you design the data warehouse, be alert for situations where denormalization may be beneficial. In general, denormalization speeds data retrieval, which may be desirable for a data warehouse. However, proceed with caution when denormalizing the data warehouse design. Denormalization is not mandatory simply because denormalized data is optimized for data access and the data warehouse is primarily “read only.” Because the data must be populated into the data warehouse at some point, a data warehouse is in actuality “read mostly,” not read only. Denormalized data is difficult to maintain and should be avoided if performance is acceptable. So, as with a traditional database implementation, denormalize only when a completely normalized design will not perform optimally.

Proceed with caution when denormalizing the data warehouse design.

Another consideration for the physical design of a data warehouse is *data compression*. Because data warehouses tend to be substantial in size, compression can be used to reduce storage requirements. The performance impact of compression will depend on the ability of your DBMS to compress data. Sometimes compression aids performance because more rows can be stored in the same amount of space and fewer I/Os are required.

Finally, the DBA should design the physical database for the data warehouse such that it encourages parallel operations. Due to the high volume of data that must be processed by the analytical queries typical of OLAP and DSS, performance is usually enhanced by parallel activities.

Consider partitioning tablespaces to take advantage of parallelism features of the DBMS, and consider basing the partitions on the usage patterns of your data warehouse access. In addition, placing indexes and data tablespaces on different disk devices can help parallel tasks to operate more efficiently.

Data Movement

A big part of any data warehousing effort is establishing the data requirements and then creating methods to capture and move the data into the warehouse. Without a clear understanding of end user needs, the DBA will have great difficulty determining what data to move, never mind how to move it.

Establishing the data requirements demands a clear understanding of end user needs.

To establish what data is required, the DBA must elicit answers to two important questions:

- What is the purpose of the data warehouse?
- What are the results expected from the data warehouse?

The next step is to ferret out the required data from the OLTP systems and databases. The data warehouse team must understand each piece of data, including its source (where it is created and by whom), its relationship to other data elements, its meaning, and what it needs (if anything) to make it viable for inclusion in the data warehouse. All of this metadata should be captured and maintained so that unanticipated changes to the data in the OLTP systems after the data warehouse goes live can be accommodated.

If a source for the data cannot be found internally, you may need to purchase data from external suppliers and integrate it into your data warehouse.

After all data has been identified, as well as any transformations that are required, processes need to be established for the extraction and movement of the data. Typically, the extraction process gathers the data from the source OLTP databases and places it in an intermediate staging area. However, this still-raw data may need to be transformed before it is placed into the data warehouse.

Some data warehouses are designed to remain in read-only mode, except

when an ETL process is occurring. Switching between read-only and read-write modes is accomplished using database commands or by setting system parameters (depending upon the DBMS being used).

Tools can be used to assist, including LOAD and UNLOAD utilities, import and export programs, propagation software, and ETL products. Be sure to include the cost of such software in your data warehousing budget, and make sure that all DBAs who need to use this software are sufficiently trained.

Consistent Data Acquisition

Businesses react and adapt to industry trends, with resulting changes to data. As the data in operational systems changes, so must the data in the warehouse. Over time, fields will be eliminated, meanings will change, sizes and types will change, and more. You must plan to keep track of physical data changes, as well as changes to the semantics of the data. Regardless of the type of change, you will need utilities, tools, and processes to keep on top of these issues and respond appropriately.

As the data in operational systems changes, so must the data in the warehouse.

Data Cleansing

The quality of the data in the data warehouse is of utmost importance to the overall usefulness of the data warehouse. If the data is not reliable, or cannot be understood by the users, the data warehouse will fail. In other words, data in the warehouse is only as good as its source. Failure to clean dirty data can result in the creation of a data outhouse instead of a data warehouse.

Therefore, the data warehouse project requires data cleansing and scrubbing tasks to be defined and performed as the data is moved from other sources to the data warehouse.

Data scrubbing can be used to “clean” data as it is moved into the data warehouse. One component of data scrubbing is the transformation of codes into meaningful values for the user. For example, a CUSTOMER-CODE of 5 means nothing to the typical user, but a CUSTOMER-CODE of “Corporation” or “Individual” is usable and helpful.

Transform codes into meaningful values for the user.

A number of insidious data problems can reduce the usefulness of a data warehouse. We've all had the experience of looking at the contents of one of our major files or databases and knowing intuitively that the data is incorrect. There is just no way that an employee was born in 1999. And that next record looks bad, too. How could she have been born in 1979 but hired in 1978?

All too often, these types of data integrity problems are glossed over. "No one would actually take that information seriously, would they?" Well, maybe people won't, but computerized systems will. Incorrect information can be summarized, aggregated, or manipulated in some way, and then populated into another data element. When erroneous data elements are moved into the data warehouse and processed analytically, your company's business can be impacted. What if warehouse data is being analyzed to overhaul hiring practices? Erroneous data could impact employment decisions if enough of those hire and birth dates were inaccurate.

Never cover up data integrity problems. Document them and bring them to the attention of your manager and the users who rely on the data. Usually, the business units using the data are authorized to modify and validate their business data.

Never cover up data integrity problems.

Small data discrepancies can become statistically irrelevant when large volumes of data are averaged. However, averaging is not the only analytical function that is employed by analytical data warehouse queries. What about sums, medians, maximums, minimums, and other aggregate and scalar functions? Even further, can you actually prove that the scope of your data problems is as small as you think it is? The answer is probably no.

This is just one small example of the data integrity violations that many application systems allow into their production data stores. Some of the integrity violations may seem to be inexcusable. For example, most of us have encountered a SEX or GENDER column (or field) that stored everything from "*" to "!" to a blank. Shouldn't it be a simple matter to programmatically force the values to be either "M" or "F"? The short answer is yes, but this simplifies that matter too much, because certain information, by law, has to be optional.

The real problem is that a systematic manner of recording "unknown" values has not been employed. Every program that can modify data should be

forced to record a special “unknown” indicator if a data value is not readily available at the time of data entry. Most relational DBMS products enable data columns to store a null to indicate “unknown” or “unavailable” information. Prerelational DBMS products and data files do not have this option. However, some specific, standard default value can be chosen. The trick is to *standardize* the default value.

Standardization of “unknown” values can be a tedious process. Our primitive example used a data element with a simple domain of two valid values. Most data elements have domains that are considerably more complex. Determining which values are valid can be difficult for someone who is not intimately familiar with the application system that allowed the values to be inserted in the first place. Is “1895-01-01” a valid date for that field, or is it a default for an “unknown” value?

Standardization of “unknown” values can be a tedious process.

Only an in-depth analysis of the programs and the metadata in the corporate repository can provide the answer. Nineteenth-century dates may be valid for birth dates, stock issuance dates, account inception dates, publication dates, and any number of other dates with long periods of “freshness.” However, just because the program allows it to be entered does not mean it is actually a valid date! It is quite simple for a user to type in 1895 instead of 1985. If the data entry program is not intelligent enough to trap these types of errors, your systems will insert dirty data into production data stores. This type of data integrity problem is the most difficult to spot. It is quite likely that only the end user most familiar with the data could spot these types of problems.

So, what is the solution? Several techniques can be used, but the best approach is to foster an environment in which data is truly treated as a corporate asset. The problem is attracting the appropriate high-level management personnel who can implement a policy that values data. But now we are moving the discussion outside the scope of data warehousing and back to data administration. So, let’s get back on track.

A key part of the data movement activities for populating a data warehouse will involve the identification of invalid and coded data, and transforming that data into useful, queryable information. Automated ETL software can help you with this, but it cannot do it all for you.

Transform the data into useful, queryable information.

Data Quality Issues

Maintaining data quality will be an ongoing concern. Both the end users and the data warehouse construction and maintenance team are responsible for promoting and fostering data quality. Data problems will be discovered not only throughout the development phase of the data warehouse but also throughout its useful life.

Be sure to establish a policy for how data anomalies are to be reported and corrected before the data warehouse is made generally available. Additionally, be sure to involve the end users in the creation and support of this policy; otherwise, it is doomed to fail. Remember, the end users understand the data better than anyone else in the organization, including the data warehouse developers and DBAs.

Establish a policy for the reporting and correction of data anomalies.

Do not underestimate the amount of time and effort that will be required to clean up dirty data. Understand the scope of the problems and the processes required to rectify those problems. Take into account the politics of your organization and the automated tools available. The more political the battle, the longer the task will take; the fewer tools available, the longer the task will take. Furthermore, if the tools you do have are not understood properly, the situation will probably be worse than having no tools at all.

Operational Problems

As the previous discussion makes clear, you will encounter problems in the operational systems feeding the data warehouse—problems that may have been running undetected for years. Your options for handling them are few:

- Ignore the problem with the understanding that it will exist in the data warehouse too.
- Fix the problem in the operational production system.
- Fix the problem during the data transformation phase of data warehouse population, if possible.

Previously undetected operational problems may be

uncovered.

Of course, the second and third options are the favored approaches. Fixing the problem in the operational system is the best solution of all, but you will need to examine the budget and deadlines associated with the data warehouse project to determine whether production problems can be remedied during the project. If not, you can fix the data problems as the data flows into the data warehouse. You should also plan on tracking the source of the errors in the production systems, if possible, so that a problem once identified is at least documented.

However, do not mix operational needs into the data warehouse project plan. When a data warehousing project is first initiated, the objectives may be a mix of operational and analytical/informational goals. This is a recipe for disaster. Redefine the project to concentrate on nonoperational, informational needs only. The primary reason for the existence of the data warehouse is to segregate operational processing from reporting.

Do not mix operational needs into the data warehouse project plan.

Data Warehouse Scalability

As a data warehouse becomes accepted within the organization, demand for its services will grow. The need for new reports and summary tables increases, and the data warehouse can explode to several times its original size. Some industry surveys indicate that 60 to 70 percent of data warehouses are filled with duplicate or redundant data such as summary tables and indexes. This can more than double the disk storage required to store the data warehouse. The more users on the system and the more simultaneous queries there are, the more potential there is for delays in response time. It is important, therefore, to architect the system so that it will be able to scale linearly with demand. The DBA can choose to use parallel processors, parallel databases, bitmapped indexes, data compression, and other techniques to provide better scalability.

Size Issues

One aspect of scalability is size. Many novices incorrectly assume that a data warehouse must be large. While it is true that many data warehouses are quite large and consume massive amounts of storage, a data warehouse can be

small and still provide value to the organization. The reality is that size doesn't matter.

The reality is that size doesn't matter.

Data Warehouse Performance

Data warehouse performance is similar to overall database performance, as discussed in [Chapter 9](#). However, different aspects of performance are emphasized. In general, data warehouse performance can be viewed from four perspectives:

- *Extract performance*—how smoothly data is updated and refined in the warehouse
- *Data management*—the quality of the data warehouse design and database maintenance
- *Query performance*—how well the database is tuned for analytical processing and how well each query is written
- *Server performance*—the configuration and performance of the hardware

A data warehouse is populated, at least in part, with data extracted from production systems. The efficiency with which data is extracted from the production system impacts the usability and performance of the data warehouse. The longer it takes to gather data and load it into the data warehouse, the longer the period that the data warehouse will not reflect current data. Moreover, that lack of currency can impact the usefulness of any data analysis performed on the data. It is therefore imperative to optimize the speed with which the data warehouse is loaded, unloaded, and accessed. High-speed database utilities can be used to optimize the flow of data throughout the life cycle of the data warehouse.

The design of the data warehouse is intrinsic to its viability. A data warehouse must be designed for easy access and analysis. The DBA must understand how the data is going to be accessed and design the database structures accordingly. With OLTP systems, dynamic activity occurs against live data. Many concurrent updates, inserts, and deletes can be going on at any time. The way in which tables are indexed is probably the most crucial element of database design for OLTP systems. In contrast, the OLAP and DSS systems used for data warehousing are read-only systems with no (or

limited) online modification capabilities—instead, data is loaded during off-peak hours. Therefore, the data warehouse may require many finely tuned indexes for daily analytical queries and fewer (perhaps no) indexes during the loads. The DBA must carefully plan to balance the efficiency of the queries against the efficiency of the load processes.

A data warehouse must be designed for easy access and analysis.

The data warehouse indexes do not have to be the same indexes that exist in the operational system, even if the data warehouse is nothing more than an exact replica or snapshot of the operational databases. The DBA must optimize the indexes, based on the access patterns and query needs of the decision support environment of the data warehouse. Additionally, consider using bitmap indexes for data warehouses to optimize the performance of your DSS applications.

Furthermore, from a database design perspective, methods for ensuring data integrity during data warehouse operations are not needed. Because data is cleansed and scrubbed during the data transformation process, implementing data integrity mechanisms such as referential integrity, triggers, and check constraints on data warehouse tables is not efficient. Even without a comprehensive cleansing, the data in the warehouse will be as good as the data in the source operational systems (which should use RI and check constraints).

Always remember that the data warehouse is a decision support vehicle. Although the data will be updated and summarized periodically, it should not be updated by the DSS users. If your design calls for a data warehouse in which all the data is modified immediately as it is changed in production, you need to rethink your data warehouse design.

The data warehouse is a decision support vehicle.

Consider starting data warehouse databases in read-only mode if your DBMS permits. Issuing a `START` command to enable read-only mode can have the additional effect of eliminating locking on the read-only databases. Of course, when the data warehouse is refreshed, the databases have to be restarted in read-write mode.

From the query performance perspective, everything is different in the data warehouse environment. Subsecond response time is not expected—indeed, a query can run all day long if the results of that query are expected to deliver significant analytical insight. That does not mean that the DBA does not need to monitor data warehouse query performance. It simply means that the goals of monitoring DSS queries are different. The DBA should examine trends instead of absolutes.

For example, in a data warehouse it would be interesting to note that queries running between 3:00 p.m. and 4:00 p.m. from the finance department on expense data average between 10 and 25 seconds response time. The DBA can then examine the daily response time by time ranges to determine whether performance is falling within expected ranges or not. If query response time suddenly and consistently jumps for certain time ranges, the DBA may need to help end users write more efficient queries. Alternatively, the DBA could change the data warehouse implementation to better support the users by building indexes or transforming data by aggregation or summarization, for example.

Because data warehouses are read only in nature, you should be able to run DSS and OLAP queries using dirty reads. Recall from [Chapter 5](#) that a dirty read is implemented with the UNCOMMITTED READ isolation level. With this isolation level, queries might read data that never actually exists in the database, because the data has been changed by another process but is not yet committed. This is rarely a concern in the data warehouse because of its read-only nature.

The server on which the data warehouse resides requires peak performance around the clock, but this demand is mitigated by the difference between OLTP and OLAP/DSS processing. However, the server hardware and configuration should not become a bottleneck for data warehouse performance and thus must be monitored and tuned.

Organizations should use an agent-based performance monitor that collects, analyzes, and stores thousands of performance measures; is configurable for multiple environments; and offers both real-time and historical perspectives on viewing all critical metrics. In this manner, organizations can implement an integrated database performance solution that is capable of monitoring and managing the performance of the following:

- Relational databases in Windows, UNIX, and mainframe environments

- Servers in distributed environments
- Entire enterprise network
- Distributed client/server transactions

Automatic Summary Tables (aka Materialized Query Tables)

Most experts agree that aggregating records in a data warehouse can palpably enhance the performance of large data warehousing applications. This is so because the vast amount of detailed information that exists in most data warehouses cannot be easily understood. Summarizing the data into logical groups makes it easier to be assimilated and analyzed in a cogent manner.

Aggregating records can palpably enhance the performance of large data warehousing applications.

The single most dramatic way to affect performance in a large data warehouse is to provide a proper set of aggregate (summary) records that coexist with the primary base records.¹

However, summarizing data poses several problems. First, summarizing data can require an enormous amount of processing. Then, once summarized, the data must be stored, which takes up valuable disk space. Even more troubling is the problem of identifying what data should be aggregated.

Many modern DBMSs provide the ability to automate aggregation using [*automatic summary tables*](#), also known as *materialized query tables* or *materialized views*. The basic concept is to create a new aggregated table from a base table by using basic SQL functions such as COUNT, SUM, AVG, MIN, or MAX. The result is a table that is aggregated at a higher-level dimension.

Regardless of the actual DBMS implementation, the basic idea is that the DBA controls the creation of summary tables. The summary table is defined by means of a query using SQL functions. Users do not need to even know that summary tables exist. The DBMS optimizer can determine when using a summary table can optimize a query written against the detail data and then automatically rewrite the query to use the summary table. Furthermore, in most cases, the DBMS automatically maintains aggregates. In other words, when underlying data changes, the DBMS automatically updates the

summary table data.

The DBA controls the creation of summary tables.

In general, automatic summary tables can significantly optimize data warehouse performance. However, the DBA will need to monitor which aggregates are used most frequently to ensure that the proper automatic summary tables are built. Additionally, the DBA will need to plan for additional disk usage to support the aggregated data.

Data Freshness

Because data is constantly changing in the OLTP systems, data warehouse processes must be set up to keep data up-to-date. As business users analyze data in the warehouse, the freshness of the data is an issue. The results of an analytical query issued against data that is up-to-date as of yesterday are likely to be more useful than the same query run against month-old data.

Data warehouse processes must be set up to keep data up-to-date.

The difference between the data in the data warehouse and the data in the production system is known as *data latency*. Business needs will drive the desired amount of latency. Some data warehouses require very fresh data (that is, low latency), whereas other data warehouses can tolerate high latency.

The DBA must establish a means for the business users to determine the freshness of the data. Whenever new or updated data is loaded into the data warehouse, the business users must be capable of ascertaining the validity date of that data. Think of it like the date that is stamped on a carton of milk.

However, do not be confused into thinking that only a completely up-to-date data warehouse is of use. When examining data for patterns and trends, the freshness of the data may be irrelevant—or, perhaps, mostly irrelevant. The importance of data freshness will vary from data element to data element, from business unit to business unit, and from query to query.

The bottom line, though, is that business users should be given an easy way to establish when the data was last loaded and how fresh the data is.

Data Content

The data warehouse administrator may need to provide metadata on the contents of the data warehouse. This metadata could include the types of data being stored, which types of data change most frequently, and how fast certain types of data are growing. Keeping track of growth patterns helps to determine the areas of the data warehouse that may require attention in terms of storage allocation, reorganization, or redesign.

Furthermore, it is possible that no matter how closely you manage the population of the data warehouse, invalid or inappropriate data will be loaded. Developing methods to analyze and verify data integrity and relationships in the data warehouse is a worthwhile endeavor.

Develop methods to analyze and verify data integrity and relationships in the data warehouse.

The primary value of the data warehouse is in its data content. Therefore, it is good practice for the DBA to gain as much understanding of that data content as possible.

Data Usage

A data warehouse management task that is not associated with traditional database administration is the identification of unused data. Maintaining certain data in the warehouse does not make sense when certain thresholds or conditions are met. This situation may occur because of usage reasons (data is no longer used by the business), technology reasons (reaching a capacity limit), regulatory reasons (change in regulations or laws), or business reasons (restructuring data or instituting different processes).

A data warehouse is usually a very large consumer of disk storage. Sometimes users go overboard when designing a data warehouse and store more data than is required. The data warehouse DBA must deploy methods to examine what data has been accessed and how frequently. The DBA should ascertain from the business users whether such inactive data is still required.

If large amounts of unused data are identified, the DBA may decide to remove the data from the warehouse. Such action should not be taken indiscriminately. Only after verifying that the data has been unused for at least a year—to take into account seasonal access variations—should any action be taken on that data.

The unused data may then be purged from the data warehouse or perhaps

moved to cheaper storage media (tape, optical disc, or CD-R, for example). Of course, purged data should be stored for a period of time to assure that it was not purged in error.

Purged data should be stored for a period of time.

The capability to track data warehouse activity is a useful administrative goal. For example, the DBA may opt to profile usage on a daily basis, and then summarize it by day of the week, week, or month. Usage by business unit or department is also a useful statistic. The more you can track data at discrete levels—even down to the column level—the better. Furthermore, this information can be coupled with performance monitoring statistics to gauge response time against what was actually accessed. In general, the more you know about who is using what data and when, the better you will be able to manage the data warehouse environment.

Financial Chargeback

In most organizations, data warehouse projects are managed by multiple departments, each of which has its own financial goals. Data warehouse DBAs should ensure that they can charge back appropriate costs to business units and users so that they can meet financial reporting requirements. An integrated solution is required that monitors IT costs by providing critical chargeback services that track information resources used organization-wide.

For chargeback purposes, you should consider measuring and tracking every request made to the data warehouse, including who requested the data, the number of rows returned by the request, the CPU cycles spent to return the data, and which tables were accessed. Depending on your particular organizational needs, you may also wish to track the specific time of day the request was made. Some organizations choose to charge different rates based on the activity in the system at specific times.

Measure and track every request made to the data warehouse.

Even if your organization chooses not to actually bill departments by usage, it can make sense to establish a financial chargeback system to demonstrate the value of the data warehouse to the business units that use it the most.

Backup and Recovery

The DBA must create a backup and recovery plan for the data warehouse even though most access is read only. Some novices mistakenly view backup and recovery as a non-issue for data warehouses, but this is simply not the case. Data is moved into the data warehouse on a consistent basis—some of it from external sources that might not be easily duplicated.

Furthermore, if warehouse data is lost due to some type of failure, the production data from which it was gathered may no longer exist in the same form. Perhaps the production database was purged to provide space for more transactions. If the data was not archived, it cannot be reproduced, and even if it was archived, it may be difficult to restore. In addition, data in the data warehouse typically exists in different forms from production data. Data may have been cleansed, transformed, or aggregated before being populated into the data warehouse. Such data may be difficult or impossible to recreate. Furthermore, data warehouses often contain external data that, if lost, may have to be purchased again, creating an unnecessary financial drain on the company.

So, for all of the reasons you need to provide a backup and recovery plan for traditional databases, you need to do the same for the data warehouse databases. Of course, the backups need be taken only whenever data is added to the data warehouse.

Indeed, backup and recovery needs special consideration within the context of the data warehouse. The data warehouse should have a backup and recovery strategy that will enable the organization to recover all essential data in an emergency. Data warehouses are becoming just as essential to businesses as transaction-oriented databases.

Backup and recovery of the data warehouse needs special consideration.

Depending on the size and nature of the data warehouse, you may decide not to back up particular pieces of the data warehouse because you can refresh the data more efficiently. Such cases, however, are more the exception than the rule. Review the cost/benefit of each data warehouse and data mart, keeping in mind how often the data is updated or refreshed and how long a recovery would take to implement.

Finally, disaster recovery requirements must not be overlooked. Organizations are becoming more and more dependent on the information

that a data warehouse provides, thereby raising the criticality of data warehouse applications. This means the data warehouse must be treated like any other critical system in terms of disaster recovery planning.

Don't Operate in a Vacuum!

The DBA must be aware of the business requirements of the data warehouse and ensure that these needs are met. This is a challenge, but it is simple when compared to the challenge of keeping the data warehouse accurate and available in light of changing operational systems.

As business needs change, operational systems change. When operational databases and files change, the data warehouse will be affected as well. When a data warehouse is involved, however, both the operational database and the data warehouse must be analyzed for the impact of changing any data formats. This is the case because the data warehouse stores historical data that you might not be able to change to a new format. Before any changes are made to the operational system, the data warehouse team must be prepared to accept the new format as input to the data warehouse and to either maintain multiple data formats for the changed data element or to implement a conversion mechanism as part of the data transformation process.

Be aware, though, that conversion can result in lost or confusing data. When undertaking a conversion, it is wise to plan thoroughly and automate the process as much as possible with tested, trusted tools and scripts. Relying solely on human conversion will almost surely result in errors in the data.

Conversion can result in lost or confusing data.

Summary

Data warehouses can provide organizations with a competitive advantage as users begin to analyze data in conjunction with business trends. Once a data warehouse has been implemented, you cannot turn back because your users will be hooked, your organization will be more profitable, and you'll have the satisfaction of contributing to the success of the business.

DBAs whose duty it is to manage their company's data warehouse need to gain an understanding of the analytical nature of data warehousing. Many of the tasks required to support a data warehouse are similar to the tasks required to support a traditional, OLTP-focused database. However, there are

differences that need to be understood and built into your data warehouse administration policies and procedures.

Managing and administering a data warehouse requires a combination of knowledgeable people, full-featured tools, and robust methods. Only through the integration of these disparate resources can the DBA assure that the data warehouse will be built and managed effectively.

Review

1. What is the difference between a data warehouse and a data mart?
2. Is it necessary to develop a backup and recovery plan for the data warehouse, since it is a read-only environment?
3. Describe a few data cleansing techniques that may need to be applied to data before it can be loaded into the data warehouse.
4. When designing a data warehouse, how should problems that are identified in the operational OLTP systems be handled?
5. Describe the star schema and why it is an effective design for data warehouse databases.
6. What are the four perspectives of data warehouse performance management?
7. When might data need to be purged from the data warehouse?
8. What is the difference between OLAP and data mining?
9. Discuss the various options that must be considered specifically when creating the physical design of the data warehouse.
10. Denormalization is always a good option for data warehouse databases: true or false?

Suggested Reading

Adamson, Christopher. *Star Schema: The Complete Reference*. New York, NY: McGraw-Hill (2010). ISBN 978-0-07-174432-4.

Barquin, Ramon, and Herb Edelstein, eds. *Planning and Designing the Data Warehouse*. Upper Saddle River, NJ: Prentice Hall (1997). ISBN 0-13-255746-0.

Devlin, Barry. *Data Warehouse from Architecture to Implementation*.

- Reading, MA: Addison-Wesley (1997). ISBN 0-201-96245-2.
- Hackney, Douglas. *Understanding and Implementing Successful Data Marts*. Reading, MA: Addison-Wesley (1997). ISBN 0-201-18380-3.
- Inmon, W. H. *Building the Data Warehouse*. 2nd ed. New York, NY: John Wiley & Sons (1996). ISBN 0-471-14161-5.
- Inmon, W. H., and Richard Hackathorn. *Using the Data Warehouse*. New York, NY: John Wiley & Sons (1994). ISBN 0-471-05966-8.
- Inmon, W. H., Claudia Imhoff, et al. *Building the Operational Data Store*. New York, NY: John Wiley & Sons (1996). ISBN 0-471-12822-8.
- Inmon, W. H., John A. Zachman, et al. *Data Stores, Data Warehousing and the Zachman Framework*. New York, NY: McGraw-Hill (1997). ISBN 0-07-031429-2.
- Kelly, Sean. *Data Warehousing in Action*. New York, NY: John Wiley & Sons (1997). ISBN 0-471-96640-1.
- Kimball, Ralph. *The Data Warehouse Toolkit: The Complete Guide to Dimensional Modeling*. 2nd ed. New York, NY: John Wiley & Sons (2002). ISBN 0-471-20024-7.
- Lewis, William J. *Data Warehousing and E-Commerce*. Upper Saddle River, NJ: Prentice Hall (2001). ISBN 0-13-091154-2.
- Poe, Vidette. *Building a Data Warehouse for Decision Support*. Upper Saddle River, NJ: Prentice Hall (1996). ISBN 0-13-371121-8.
- Turban, Efraim. *Decision Support and Expert Systems*. Englewood Cliffs, NJ: Prentice Hall (1995). ISBN 0-02-421701-8.
- Witten, Ian H., and Eibe Frank. *Data Mining: Practical Machine Learning Tools and Techniques*. San Francisco, CA: Morgan Kaufmann (2005). ISBN 0-12-088407-0.

21. Database Connectivity

Installing the DBMS and creating databases are vitally important parts of the DBA's job. However, if the DBMS is not connected to the network, no one will ever be able to access the data.

Multitier, Distributed Computing

This section offers a short introduction to multitier, distributed computing, including client/server. The DBA must possess an understanding of the terms commonly used in a multitier processing setting to assure database access in a distributed business environment.

A Historical Look

In the early days of computing, computers required large amounts of power and resources to run. Batch processing was the only operating mode, and users received printed reports as their only source of information. Mainframe programmers communicated with computers through punched cards. As technology progressed in the 1970s, online processing gave the user more power and closer availability to data. In the 1980s and 1990s, personal computers became widely available and affordable. Inexpensive software accelerated the growth of the PC, which today is a machine with considerable power and memory capacity. Some are as powerful as the older mainframes.

Initially, PCs were used as stand-alone equipment with their own resident data and software. The only way users could share data was by copying it to a diskette and giving it to someone else. PCs were also used to emulate mainframe terminals. Although terminal emulation helped somewhat, it did not add significantly to the overall productivity.

PCs were initially used as stand-alone equipment.

As PC usage expanded, end users became more computer literate. This initiated an insatiable desire for more data and more connectivity to corporate systems. Seamless access to corporate databases on many platforms became a requirement. However, for mission-critical applications to achieve productivity enhancements, the computing assets—mainframe, midrange, and personal—of the organization had to be integrated. Deploying the correct

computing resources to resolve the issues for which they were best suited optimized productivity. Over time, a company's overall competitiveness could be impacted if end users had to spend an inordinate amount of time and effort simply to access and analyze corporate data—wherever it resided.

The need to share information gave way to networks of workgroups using the same data, printers, and software. This helped IT departments to include PCs in traditional system management tasks such as backup and technical support. As users grew more accustomed to PCs and their simpler interface, expectations for mainframe and midrange applications grew. The use of icons, dropdown menus, and multiple colors became a requirement of new applications that could be accommodated easily on a high-powered personal computer but were not commonplace for corporate mainframe applications.

Organizations have adopted and integrated technology across multiple computing platforms. Today, distributed, multitier processing is commonplace and the network keeps these disparate platforms connected. IT personnel need to administer a complex environment to deliver computing services to the business.

Distributed, multitier processing is commonplace today.

With the advent of the World Wide Web, users became more comfortable interacting with their computers through a Web browser. Today, the Web browser is the most popular client through which applications are accessed.

Business Issues

Three forces in business have driven the change to multitier environments: downsizing, upsizing, and rightsizing. As the name implies, *downsizing* is a transition downward, from mainframes to networked PC workstations. The proliferation of more powerful, less expensive PCs has made downsizing an attractive option for many large mainframe shops. *Upsizing* is the move from stand-alone PCs to a network that connects many PCs together and provides access to corporate databases. Finally, *rightsizing* is the practice of placing applications on the most appropriate platform for the requirements. For example, applications requiring high availability and scalability may require a mainframe, at least for some pieces of the application. In contrast, an application requiring rapid delivery and a highly customized interface will likely require a PC workstation running some version of Windows for some

of its components.

What Is Client/Server Computing?

Client/server computing enables a user on a PC to run standard business applications and access corporate shared resources over a network connection. Basically, *client/server* can be defined as a relationship between two or more computing entities that function together to accomplish a task. Each of these entities performs a portion of the work involved. One entity (the “client”) requests services and the other (the “server”) fulfills the requests. Of course, a client/server environment may use multiple servers, each fulfilling different types of requests—database access, application processing, file access, printing, and so on.

Client/server can be defined as a relationship between two or more computing entities that function together to accomplish a task.

Today Client/Server Computing Is Ubiquitous

Although the term *client/server* is not used as commonly as it was in the 1990s, client/server computing still forms the basic infrastructure upon which most modern applications are run.

Various interconnected servers enable common computer resources to be shared in a client/server implementation. Most PC-based business software applications access these networked files and printers as if they were directly connected to the desktop PC. For example, in a networked environment, the end user perceives a common pool of disk drives, printers, and other resources without necessarily being able to distinguish between local and remote resources. The most commonly used types of servers include print servers, file servers, database servers, and application servers.

Print servers provide network users access to multiple printers as a sharable resource. Instead of having a printer connected to each desktop PC, print servers provide easy access to a printer by any PC connected on a network.

File servers provide a means of sharing data files among users on the network. A file server is not intended to perform computational tasks, nor

does it run programs on behalf of clients. Its primary purpose is simply to enable the storage and retrieval of data.

Database servers advance the data-sharing capabilities of the client/server system by enabling targeted data access. Unlike LAN-based file servers, which require the entire file to be retrieved, a database server is able to accept SQL queries and respond with only the requested data.

Application servers provide services for business applications. One advantage of application servers is their ability to reuse common business logic.

When a client/server system is implemented, the designers must determine which parts of the application should run on which computing assets. A client/server application consists of the following three software layers:

- *Presentation logic* consists of the tasks required to display information on a computer screen. For client/server systems, this usually involves a GUI with point-and-click, drag-and-drop capabilities, interactive help facilities, and other ease-of-use functionality.
- *Business logic* delivers the core elements of an application needed by the end user to manipulate information for managing business. This business logic is unique to the particular business at hand and incorporates methods for implementing business strategies, conducting business transactions, and enforcing company policies.
- *Database management systems* are used by most client/server systems to provide quick access to structured data in a secure manner and to facilitate data modification while enforcing data integrity. File storage systems store the data managed by the DBMS (and handle data that is not managed by a DBMS).

The presentation-logic layer is the screen (or report) that the user sees—the graphical interface. The application-logic layer provides computer processing that allows business tasks to be accomplished. Underneath both of these software layers is the data-management layer. If a business task requires information from multiple databases, the data access facilities should be consistent and transparent to the front-end software components.

The presentation-logic layer is the screen that the user sees.

In a client/server environment, each of these components, though

independent, interoperates with the other components to provide the end user with a seamless end-to-end application. The end user does not need to know that separate components are being accessed over a network to deliver application functionality. Client/server processing allows the physical placement of two or more components on separate computing platforms and systems, as shown in [Figure 21.1](#).

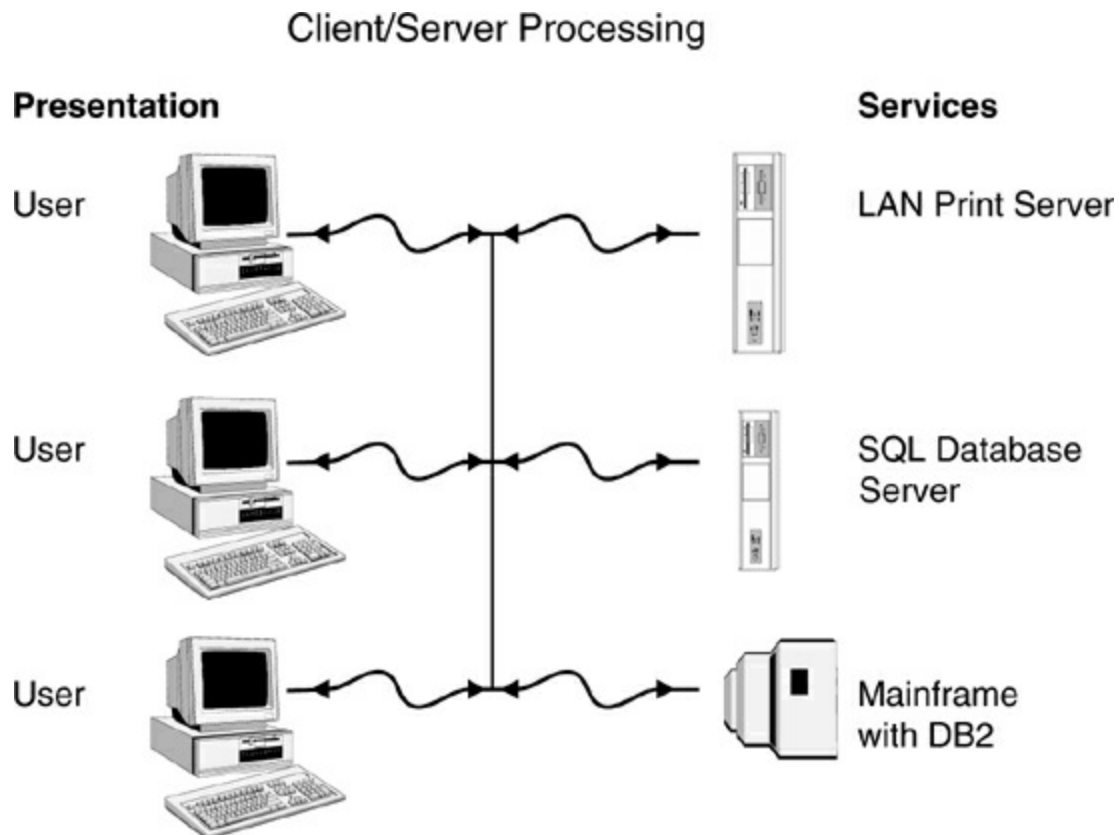


Figure 21.1 Client/server processing

The *client* computer typically is a PC connected to a network. The client manages the user presentation and provides access to the network resources. When a network service is needed, the client sends a request over the network to the desired server.

A *server* typically consists of powerful hardware and the networking components necessary to connect to clients, as well as the software to fulfill service requests made by clients. Servers must be powerful enough to handle multiple requests from multiple clients. Timely response from a server is necessary to allow client applications to appear as an integrated and unified application. The hardware used for the server can be a powerful PC (perhaps

with multiple CPUs), a UNIX-based machine, a midrange computer, or a mainframe computer. [Table 21.1](#) outlines the types of computing platforms that are most appropriate for the different client/server component layers.

Timely response from a server is necessary.

From a data management perspective, the database server provides clients with access to shared data. The server must control that access as well as activities such as data backup and recovery, performance management, and error handling.

Of course, solid networking capabilities are required to create a client/server system. The network is the enabling technology for client/server systems and much of modern computing systems. Without networking capabilities, computerized systems would not be able to connect to one another.

Centralized versus Distributed Processing

[*Centralized processing*](#) uses a host computer that is accessed via dumb terminals. PCs can be connected to a centralized system, but their only function in a centralized system is basic terminal emulation. The host-based system cannot access data stored on the PC, and any PC-based applications cannot access host data for local desktop processing (without first configuring software to enable PC-to-host connectivity). With centralized systems, the PC community is isolated from integrated, seamless processing.

In centralized processing, all data is stored on a large, centralized computer, typically a mainframe. All users run predefined applications, sometimes referred to as legacy applications, that have been built to access that data. Maintenance and porting of centralized applications to non-mainframe computer technology is expensive and difficult.

In centralized processing, all data is stored on a large, centralized computer.

Table 21.1. Recommended Hardware for Client/Server Component Layers

Component	PC	Midrange	Mainframe
User presentation	Excellent	Poor	Poor
Business logic	It depends	Generally good	Generally good
Database management	Poor	Very good	Excellent

[Distributed processing](#) involves the utilization of multiple computers to fulfill a service request. As the name implies, business tasks are distributed or allocated to multiple processors. The goal is to facilitate processing on the computing platform that is most efficient. A network that links together multiple computers is the first requirement for distributed processing. Additionally, system software must be capable of allocating tasks to different computers.

[Distributed processing](#) uses multiple computers to fulfill a service request.

Another term that is sometimes associated with distributed applications is *cooperative processing*. This refers to the functionality of two nodes *cooperating* and sharing the processing duties for a business problem. For example, one node may run some basic verification routines for new data entered, while the other node performs more rigorous testing of business rules.

Types of Client/Server Applications

In a client/server system, applications are modularized to allow them to execute over two or more nodes. The client application makes requests to a server node, which may be filled by different servers, depending on the type of request. By using software modules, client/server processing allows the sharing of commonly used software and other requested resources.

Client/server networks can be set up in a wide variety of configurations, depending on the needs of the application and the computing resources available to the organization. Keep in mind that a client/server system normally consists of three software layers: the presentation logic, the application or business logic, and the DBMS. How the processing of these layers is distributed will greatly influence the design, implementation, ease of use, and overall response time. These tasks can be broken apart in various ways across the network as shown in [Figure 21.2](#). The first and last lines in

this figure are not client/server models, because all of the processing is centralized—either on a PC workstation or a high-end server. Let’s examine a few of the client/server configurations.

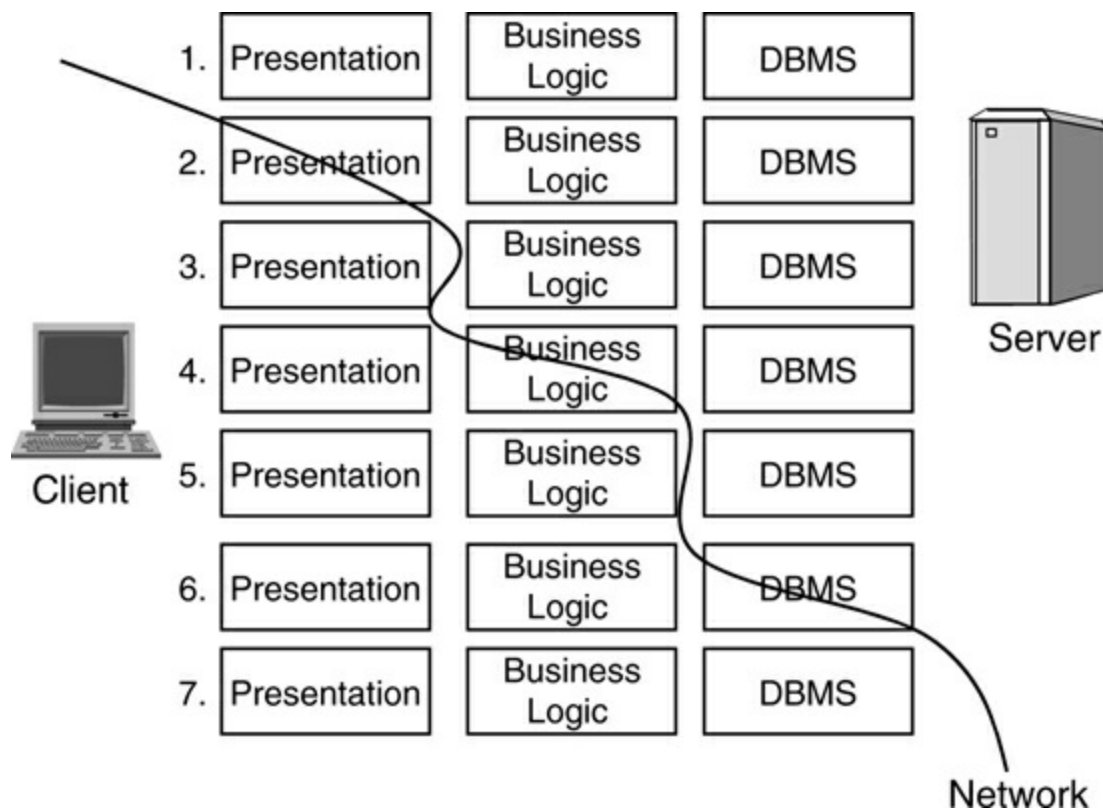


Figure 21.2 Distributing tasks across the network

Client/server networks can be set up in a wide variety of configurations.

Line 2 in the figure shows a common client/server variation: the *distributed user presentation model*. In this model, the user presentation logic is split between the workstation client and the server. The client performs the graphical and menu-driven interface elements and handles some tasks such as physically updating the screens and the rules for how the interface looks and acts. Other presentation-level tasks such as data validation may be performed on the client or on the server. Nevertheless, to the end user, the presentation layer appears to be handled entirely by the workstation. Business logic and data management tasks are performed on the server.

Line 3 in the figure shows the *decentralized user presentation model*, where all of the presentation logic is performed on the client workstation.

Dividing the processing in this manner works well, especially when client/server technology is to be phased into an organization slowly over time. Ideal applications for this model are ones that require some restructuring, but rewriting is not justified by the business needs.

Line 4 shows an alternative form of decentralization where portions of the business logic are implemented on the client and other portions are implemented on the server. This model is difficult to implement and maintain because application code that controls business functionality resides on multiple platforms.

Line 5, the *centralized data management model*, places all presentation and application logic for business tasks on the client and all data access and management software on the server. A DBMS database server handles all data requests and modifications. This provides a secure environment to administer and control corporate information. It also helps to ensure a high degree of data integrity.

Line 6 shows an alternative approach, the *distributed data management model*, where data resides both on the client node and on the server. The advantage to this approach is that queries to the local DBMS are faster. From a maintenance perspective, the administration required for a large number of distributed DBMSs is more complex. Consider the way in which ATMs work. If you use your own bank's ATM, the response time typically is very fast. This is so because the transaction accesses information on your account locally on the bank's server (probably a mainframe computer). However, if you use another bank's ATM, the response time is often much slower because the transaction goes through a central clearinghouse. In other words, information on bank accounts is distributed among all the banks connected to the clearinghouse.

Line 7, as mentioned before, is not a client/server system; it is a centralized system on a client machine.

As a DBA, you will be required to assist client/server application developers to help determine the best location for data in terms of efficiency, ease of management, and the capabilities of your DBMS.

Of course, not every client/server model is adequately described by the simple diagram we have used thus far. Indeed, a popular client/server model in use today is the *multitier implementation*, as shown in [Figure 21.3](#).




Figure 21.3 Multitier client/server implementation

One of the most popular client/server models in use today is the multitier implementation.

In this client/server implementation, two servers are deployed: one for the application and another for the DBMS. The client handles all (or most) presentation-logic duties. This is the model deployed by many popular third-party applications such as SAP ERP and Oracle's PeopleSoft and Siebel applications. Supporting such applications requires networking expertise as well as a high degree of knowledge that spans multiple platforms (because each tier is implemented using a different type of computer).

Thin Client or Fat Client

Distributed and decentralized user presentation client/server implementations are sometimes referred to as *thin client* because only limited functionality resides on the client computer. Client/server models where most of the processing occurs on the client are called *fat client*. Referring back to [Figure](#)

[21.2](#), the lines toward the bottom of the diagram (in particular, lines 5 and 6) represent fat client implementations because more logic resides on the client than on the server.

Modern multitier distributed applications prefer thin client to fat client implementations because it is easier to manage and upgrade thin client. Distributing changes to hundreds, or perhaps thousands, of nodes can be difficult when the program logic resides on multiple clients throughout an organization.

Modern client/server applications prefer thin client to fat client.

Indeed, the predominant client interface for most modern applications is the Web browser.

Network Traffic

In general, network traffic is the primary source of poor performance in client/server and multitier applications. The less data that is sent across the network, the better the performance will be. The DBA should look for opportunities to decrease the size of requests sent across the network both from the client to the server and from the server back to the client. Stored procedures can be a useful tactic because issuing a call to a stored procedure over the network usually will require less bandwidth than issuing several SQL statements. Furthermore, the stored procedure, if designed properly, can execute multiple requests on the server and return to the client only the data that is required. This helps to reduce network traffic.

Network traffic is the primary source of poor performance in client/server applications.

Additional methods can be used to minimize network traffic. For example, design the application programs to minimize the number of individual network requests. Perhaps several requests can be grouped together or combined to reduce network calls.

The DBA should investigate configuration options for the data distribution methods used by the DBMS, as well as database gateway options for reducing network traffic.

Database Gateways

In order for an application to access a database over a network, [*database gateway*](#) software is required. The database gateway allows clients to access data stored on database servers. The major DBMS vendors provide gateway software for their DBMS products, and several third-party vendors provide offerings that are more heterogeneous.

When selecting a database gateway, you will need to know what types of databases need to be accessed and what networking protocol will be used. For example, if DB2 and Oracle databases must be accessed, the gateway will need to support both of these DBMS products. Similarly, the gateway may need to support multiple networking protocols if, for instance, a combination of TCP/IP and SNA networking is used. Example database gateway offerings include IBM's DB2 Connect, DataDirect's Shadow, and Oracle Database Gateways.

The gateway may need to support multiple networking protocols.

The DBA sets up and configures the database gateway software, usually with the assistance of a network specialist. When setting up the database gateway, you will need to decide whether the gateway should be installed on every client machine or a server should be set up to handle connection requests. If the gateway resides on each client, coordinating software distribution during an upgrade of the database gateway software can be difficult. Implementing the database gateway as a server simplifies software upgrades but requires connection requests to be routed to the connection server.

Furthermore, the configuration process may require the DBA to set DBMS configuration parameters or update system catalog tables to specify the locations and protocols used by the database gateway. The proper installation and configuration of the database gateway often can be one of the most demanding and difficult DBA tasks.

Some of the features to look for in a database gateway include

- Location transparency, meaning that users do not need to be aware of the physical location of data, just the names of the tables to be accessed
- Commit transparency, meaning that distributed and local transactions

are committed using standard SQL COMMIT

- Two-phase COMMIT, meaning that distributed transactions commit data appropriately across the heterogeneous database systems
- Automatic data type translation between the disparate DBMS products in use
- Both query and update capabilities for all the heterogeneous DBMS products in use

Database Drivers

Applications can access databases using a database driver, which is a software component that uses a standard API (ODBC or JDBC). The primary functionality of a database driver is to process API calls, submit SQL requests to the database, and return results to the application.

A database driver requests and retrieves data over the network from the database server using database protocol packets, typically delivered using TCP/IP. Although core database driver functionality sounds simple enough, performance can vary greatly. It is possible to improve (or degrade) application performance significantly based on nothing other than changing the database driver.

The database driver translates standards-based API calls into low-level database requests. Of course, database drivers perform other tasks such as data conversion, buffering, error handling and mapping, load balancing, and connection pooling, to name a few. The manner in which a database driver communicates with the database depends on its architecture. Some database drivers communicate directly to the database server, whereas others require translation from one protocol to another. The manner in which the database driver performs all of these functions can have a significant impact on the performance of your database applications.

Types of Database Drivers

The first database drivers were ODBC drivers. Recall from [Chapter 5](#) that ODBC (Open Database Connectivity) is a standard C programming language interface for accessing data from a database server. In theory, ODBC is independent and can work across multiple DBMSs and operating systems. This requires an ODBC driver to translate requests from the application to the database server. The application uses ODBC functions through an ODBC

driver manager with which it is linked, and the driver passes the query to the DBMS. There are numerous providers of ODBC database driver software.

With the burgeoning popularity of Java, a similar CLI-based technique was needed to communicate between Java programs and database systems. Hence the introduction of JDBC (Java Database Connectivity). A JDBC driver provides the same type of functionality as an ODBC driver, but there are four different types of JDBC database drivers, distinguishable by their architecture.

The *bridge architecture* is the first type of database driver. This type of database driver, also known as a Type 1 driver, bridges from one standard to another. A Type 1 JDBC driver accepts JDBC calls and bridges them to ODBC calls, which are then sent to the database. The Type 1 driver is platform dependent using ODBC, which in turn depends on native libraries of the underlying operating system on which the Java Virtual Machine (JVM) is running. A Type 1 driver typically is not optimal in terms of performance and security and, worse, may not support newer functionality.

The second type of database driver is *client based*. This type of database driver, also known as a native or Type 2 driver, communicates with the database through the client software of the DBMS. In other words, a Type 2 driver is proprietary. Perhaps the biggest drawback of the Type 2 driver is the overhead of managing the installation, configuration, and maintenance of software on every client machine. When you install a new version of the DBMS, you generally also must install a new version of the Type 2 driver on all client machines.

The third type of database driver uses a *network protocol architecture*, also known as a Pure Java or Type 3 driver. The Type 3 driver makes use of a middle tier between the calling program and the database that converts calls into the vendor-specific database protocol. The Type 3 driver offers better performance than Type 1 and Type 2 drivers because it reduces network bandwidth requirements and eliminates latency. Because the communication between client and server is database independent, there is no need for vendor software to be installed on the client machine. Also the client-to-middleware software need not be updated for a new DBMS version. The primary drawback is that database-specific coding is required in the middle tier.

The fourth type of database driver communicates directly with the DBMS, eliminating the need for client software. The *database wire* driver, better

known as a Type 4 driver, offers the best performance and maintainability of the various types of database drivers. Type 4 drivers install inside the JVM of the client. This provides better performance than Type 1 and Type 2 drivers as it does not have the overhead of conversion of calls into ODBC or database API calls. And unlike Type 3 drivers, it does not need associated software to work. The Type 4 driver can immediately connect to the database without configuring additional software. Because there is no client software to install, configure, and maintain, the DBA workload to support this type of driver is minimal.

Guidance

In general, you should work to minimize the amount of data that is communicated between the database driver and the database server. The more you can limit network traffic of this nature, the better your applications should perform.

Connection Pooling

Database applications that require network connection can use connection pooling to improve performance. With connection pooling enabled, when the application requests a connection, it can obtain it from a pool of previously used connections. When the connection is no longer required, instead of the connection shutting down, it gets saved in a connection pool.

A connection pool is simply a cache of database connections maintained so that the connections can be reused when future requests to the database are required. When connections are reused, the overhead of creating a connection each time one is required can be eliminated. This requires no change to the application code.

Connection pooling can be implemented by an application server or within the database driver. Parameters are required to set the initial, minimum, and maximum sizes of the connection pool, as well as to configure how long connections should remain in the pool for reuse.

Databases, the Internet, and the Web

Connecting databases to the Web is not as difficult as managing the databases once they have been connected. As with any server, the only way to enable access to a database server over the Web is to enable some form of telecommunications connection between the database server and the Internet.

This can be accomplished in many ways, from the slow and simple approach of using a dial-up connection to permanent and faster connections using frame relay, ISDN, or T1 and T3 lines. Of course, you also may choose to use an Internet service provider (ISP) to host your Web site and the databases used by your Web applications.

Many larger organizations choose to host their own databases instead of turning them over to an ISP. By hosting your own databases, you have a better degree of control over the data—the data resides on site, is accessible by your organization even if the Internet connection goes down, and is more easily secured. When an ISP hosts your database, the data is under the control of an external entity, not your organization. This might be acceptable for smaller databases, noncritical data, or very small organizations lacking the capacity to manage and administer databases in the DBMS environment.

By hosting your own databases, you have better control over the data.

Some service providers simply set up and enable the server hardware, while you install, configure, and maintain the software (DBMS) that is run on that server. This may be a suitable compromise to turning over all control to the ISP. Choosing to host your database server with your ISP can dramatically reduce costs because dedicated connections can be quite expensive.

When you host database servers that are designed to serve Internet users, following these guidelines can help:

- Favor dedicated hardware for the database server. Do not mix and match Web server software and database server software on the same computer. Dedicated hardware makes it easier to separate database and Web resources for tuning and management.
- Favor using the fastest network connections available for connecting the application logic to the database server and the database server to the Internet connection. The faster the network connection is, the better the performance will be.

Internet-Connected Databases

Organizations are rapidly adopting Internet technologies and integrating them into their business processes—in other words, businesses are becoming e-

businesses. However, the duties of the DBA become more difficult when it comes to managing databases that are connected to the Internet. The first factor complicating Web database administration is the increased reliance on availability.

Availability

Because an e-business is an online business, it never closes. There is no such thing as a batch window for an e-business application. Customers expect full functionality on the Web regardless of the time of day. And remember, the Web is worldwide—when it is midnight in Chicago, it is 3:00 P.M. in Sydney, Australia. An e-business must be available and operational 24 hours a day, 7 days a week, 366 days a year (do not forget leap years). It must be prepared to engage with customers at any time or risk losing business to a company whose Web site is more accessible. Some studies show that if Web users click the mouse and do not receive a transmission back to their browser within 7 seconds, they will abandon that request and go somewhere else.

An e-business never closes.

The net result is that e-businesses are more connected, and therefore must be more available in order to be useful. As e-businesses integrate their Web presence with traditional IT services such as database management systems, they are creating heightened expectations for data availability. In fact, the term *e-availability* is sometimes used to describe the increased availability requirements of Web-enabled databases.

E-availability describes the level of availability necessary to keep an e-business continuously operational. Downtime and outages are the enemies of e-availability. Recall from the discussion in [Chapter 8](#) that 70 percent of application downtime is caused by planned outages to the system—and only 30 percent by unplanned outages.

Downtime and outages are the enemies of e-availability.

But what does all of this mean for the e-DBA? Because e-DBAs are charged with maintaining the high degree of availability required by e-business, one of their first priorities is to minimize downtime resulting from planned outages. How can an e-DBA accomplish this? The best way to reduce downtime is to avoid it. The following steps can be taken to avoid or

reduce downtime:

- Perform routine maintenance while systems remain operational.
- Exploit the features of the DBMS that promote availability.
- Automate DBA functions.
- Consider high-speed utilities and third-party automation solutions.
- Exploit hardware technologies.

These techniques are covered in detail in [Chapter 8](#). Nevertheless, sometimes downtime cannot be avoided. If this is the case, performing tasks faster can shorten the duration of an outage. Be sure that you are using the fastest and least error-prone technology and methods available to you. For example, if a third-party RECOVER, LOAD, or REORG utility can be run in half or a quarter of the time of a traditional database utility, consider migrating to the faster technology. In many cases, the faster technology will pay for itself much quicker in an e-business than in a traditional business because of the increased availability requirements.

Performing tasks faster can shorten the duration of an outage.

Another way to minimize downtime is to automate routine maintenance tasks. For example, recall from [Chapter 7](#) how changing database structures can be a difficult task. Tools are available that enable DBAs to make any desired change to a relational database using a simple online interface. Pointing, clicking, and selecting with the tool generates scripts that automate the required database changes. When errors are avoided using automation, downtime is diminished, resulting in greater e-availability. [Chapter 7](#) provides extensive coverage of database change management issues.

Database Design

When databases are designed for e-business applications, the tendency is for the DBA to get swept up in the dynamics of Web-based design and development. This can be a dangerous mistake when designing databases. Let's investigate the impact of e-business on the design process and discuss the basics of assuring proper database design.

One of the biggest problems that a DBA will encounter when moving from traditional development to e-business development is coping with the mad rush to "Get it done NOW!" Industry pundits have coined the phrase

“Internet time” to describe this phenomenon. Basically, when a business starts operating on “Internet time,” things move faster. One “Web month” is said to be equivalent to about three standard months. The nugget of truth here is that Web projects move very fast for a couple of reasons:

- Business executives want to conduct more and more business over the Web to save costs and to connect better with their clients.
- Everyone else is moving fast, so you had better move fast, too, or risk losing business.

When a business operates on “Internet time,” things move faster.

Rapid application development (RAD) techniques have been around for about three decades now and have been used with varying degrees of success. Sometimes RAD is required for certain projects. But RAD can be bad for database design. Why? Applications are temporary, but data is permanent. Organizations are forever coding and recoding their applications—sometimes developing the next incarnation of an application before releasing the last one.

But when did you ever throw away data? Oh, sure, you may redesign a database or move from one DBMS to another. But what did you do with the data? Chances are, you saved the data and migrated it from the old database to the new one. Some changes had to be made, maybe some external data was purchased to combine with the existing data, but the data lives on. To glean the most value from your data, it is wise to take care when designing the database. It is easier to navigate a well-designed database and therefore easier to retrieve meaningful data from it. [Chapters 3](#) and [4](#) outline the techniques necessary for designing and implementing a well-constructed database.

But what if you are forced to design a database on “Internet time”?

Well, the best advice I can give you is to be aware of design failures that can result in a hostile database. A *hostile database* is a database that is difficult to understand, hard to query, and takes an enormous amount of effort to change. Of course, it is impossible to list every type of database design flaw that could be introduced to create a hostile database. However, let’s examine some common database design failures.

Be aware of design failures that can result in a hostile

database.

Assigning inappropriate table and column names is a common design error made by novices. Database names that are used to store data should be as descriptive as possible to allow the tables and columns to document themselves, at least to some extent. Application programmers are notorious for creating database naming problems, such as using screen-variable names for columns or coded jumbles of letters and numbers for table names.

When rushed for time, some DBAs resort to designing the database with output in mind. This can lead to flaws such as storing numbers in character columns because leading zeroes need to be displayed on reports. This is usually a bad idea for a relational database. If the column is created as a character column, the developer will need to program edit checks to validate that only numeric data is stored in the column. Yet, as discussed earlier, it is better for integrity and efficiency to let the database system perform the edit checking. Therefore, data storage should be based on its domain. Users and programmers can format the data for display, instead of forcing the data into display mode for database storage.

Don't design the database with output in mind.

Another common database design problem is overstuffing columns. This actually is a normalization issue. Sometimes, for convenience, a single column is used to store what should be two or three columns. Such design flaws are introduced when the DBA does not analyze the data for patterns and relationships. An example of overstuffing would be storing a person's name in a single column instead of capturing first name, middle initial, and last name as individual columns.

Poorly designed keys can wreck the usability of a database. A primary key should be nonvolatile, because changing the value of the primary key can be very expensive. When you change a primary key value, you have to ripple through foreign keys to cascade the changes into the child table. A common example is using a Social Security number for the primary key of a personnel or customer table. This is a flaw for several reasons. First, a Social Security number is not necessarily unique. Second, Social Security numbers are unique to the United States, so if your business expands outside the USA, what do you store as the primary key?

Poorly designed keys can wreck the usability of a database.

Actually, failing to account for international issues can have greater repercussions. For example, when storing addresses, how do you define zip codes or postal codes? Many countries use postal codes that are not necessarily numeric. And how do you define state or province? How do you assure that you capture all of the information for every person to be stored in the table, regardless of country? The answer, of course, is proper data modeling and database design.

Denormalization of the physical database is a design option, but it can be done only if the design was normalized to begin with. Actually, a more fundamental problem with database design is improper normalization. By focusing on normalization, data modeling, and database design, you can avoid creating a hostile database.

Without proper up-front analysis and design, the database is unlikely to be flexible enough to easily support the changing requirements of the user. With sufficient preparation, flexibility can be designed into the database to support the user's anticipated changes. Of course, if time is not taken during the design phase to ask the users about their anticipated future needs, you cannot create the database with those needs in mind.

If data is the heart of today's modern e-business, database design is the armor that protects that heart. Data modeling and database design are the most important parts of creating a database application. If proper design is not a component of the database creation process, you will wind up with a confusing mess of a database that may work fine for the first application, but not for subsequent applications. And heaven help the developer or DBA who has to make changes to the database or application because of changing business requirements. That DBA will have to try to tame the hostile database!

Database design is the armor that protects the heart of e-business.

Web Development and Web Services

Another popular type of application development technique relies upon Web services. A Web service is a software system designed to support interoperable machine-to-machine interaction over a network. Other services

interact with a Web service to conduct business. Typically, the Web service is engaged using service-oriented architecture (SOA) and a SOAP message.

What Is SOAP?

SOAP, or Simple Object Access Protocol, is a protocol specification for exchanging structured information in a Web services implementation. It relies upon XML for its message format, as well as other protocols such as HTTP and SMTP.

What Is REST?

The name “representational state transfer” purposely builds the image of how a well-designed Web application should behave. The user progresses through the application by selecting links (state transitions) via a network of Web pages, where each link results in the next page (representing the next state of the application) being transferred to the user.

At any rate, Web services are a form of multitier processing but designed using components that deliver specific functionality, or a service. Instead of defining an API, SOA defines the interface in terms of protocols and functionality. SOA separates functions into distinct units, or services, that developers make accessible over a network in order to allow users to combine and reuse them in the production of applications.

Another technique is to use representational state transfer, or REST. REST-style architectures are based on the client/server model. Requests and responses are built around the transfer of representations of resources. A resource can be any coherent and meaningful concept that may be addressed. A resource is usually a document that captures the current or intended state of a resource.

The client sends requests when it is ready to make the transition to a new state. The client is considered to be in transition as long as one or more requests are outstanding. The representation of each application state contains links that may be used the next time the client chooses to initiate a new state transition.

A RESTful Web service is implemented using HTTP and the principles of

REST. SOAP, on the other hand, encourages each application designer to define a new vocabulary of nouns and verbs, disregarding many of HTTP's existing capabilities such as authentication, caching, and content type negotiation. Some believe this aspect of SOAP causes too much reinvention.

Summary

Modern databases are designed and implemented for use across a network where users are connected to the database from many different types of computing platforms. The way in which databases are designed, administered, and supported will impact the usability of connected databases. A different set of issues and problems must be tackled for client/server and Internet-enabled databases. DBAs should be prepared to handle these problems by becoming well versed in not only how to connect databases but also how to manage those databases in an efficient manner once they are connected.

Review

1. Name the four types of database drivers, and discuss the pros and cons of each type.
2. Why is database design a potential problem area for Internet-enabled database application development projects?
3. What is the difference between a fat client and a thin client?
4. In a client/server environment, there are always multiple clients communicating with only one server: true or false?
5. A client/server application comprises what three components?
6. What is the biggest DBA challenge when managing an Internet-enabled database?
7. What is connection pooling?
8. Cite examples of flaws introduced during the design phase of an Internet-enabled database that could lead to the creation of a hostile database.
9. Your organization deploys PCs, midrange systems, and mainframe computers. Discuss the three layers of a client/server application and which computing platform would work best for each layer.

10. What is e-availability?

Suggested Reading

Abiteboul, Serge, et al. *Data on the Web: From Relations to Semistructured Data and XML*. San Francisco, CA: Morgan Kaufmann (2000). ISBN 1-55860-622-X.

Appelquist, Daniel K. *XML and SQL: Developing Web Applications*. Boston, MA: Addison-Wesley (2002). ISBN 0-201-65796-1.

Ceri, Stefano, et al. *Designing Data-Intensive Web Applications*. San Francisco, CA: Morgan Kaufmann (2003). ISBN 1-55860-843-5.

Feiler, Jesse. *Database-Driven Web Sites*. San Francisco, CA: Morgan Kaufmann (1999). ISBN 0-12-251336-3.

Gagliardi, Gary. [*Client/Server Computing*](#). Englewood Cliffs, NJ: Prentice Hall (1994). ISBN 0-13-290784-4.

Goodson, John, and Robert A. Steward. *The Data Access Handbook*. Upper Saddle River, NJ: Prentice Hall (2009). ISBN 0-13-714393-1.

Hackathorn, Richard D. *Enterprise Database Connectivity*. New York, NY: John Wiley & Sons (1993). ISBN 0-471-57802-9.

Lewis, Philip M., Arthur Bernstein, and Michael Kifer. *Databases and Transaction Processing*. Boston, MA: Addison-Wesley (2002). ISBN 0-201-70872-8.

Orfali, Robert, Dan Harkey, and Jeri Edwards. *Essential Client/Server Survival Guide*. New York, NY: Van Nostrand Reinhold (1994). ISBN 0-442-01941-6.

Purba, Sanjiv, ed. *High-Performance Web Databases*. Boca Raton, FL: Auerbach (2000). ISBN 0-8493-0882-8.

Riccardi, Greg. *Principles of Database Systems with Internet and Java Applications*. Boston, MA: Addison-Wesley (2001). ISBN 0-201-61247-X.

Salemi, Joe. *Guide to Client/Server Databases*. Emeryville, CA: ZD Press (1993). ISBN 1-56276-070-X.

Sheldon, Tom, ed. *LAN Times Guide to Interoperability*. Berkeley, CA: McGraw-Hill (1994). ISBN 0-07-882043-X.

Terplan, Kornel, and Jill Huntington-Lee. *Applications for Distributed Systems and Network Management*. New York, NY: Van Nostrand Reinhold (1995). ISBN 0-442-01873-8.

Therriault, Marlene L. *Oracle8i Networking 101*. Berkeley, CA: McGraw-Hill/Oracle Press (2000). ISBN 0-07-212517-9.

Toledo, Hugo, Jr. *Oracle Networking*. Berkeley, CA: McGraw-Hill/Oracle Press (1996). ISBN 0-07-882165-7.

22. Metadata Management

In addition to managing data, DBAs need to be able to manage and control the definition of the data elements used in databases. Without an understanding of the structure, limitations, definition, and description of data, it is likely that data will be misinterpreted or misused. Furthermore, data that is not well defined can cause database integrity problems.

What Is Metadata?

Have you ever watched the *Antiques Roadshow*¹ program on television? In this show, people bring items to professional antique dealers to have them examined and evaluated. The participants hope to learn that their items are long-lost treasures of immense value. The antique dealers always spend a lot of time talking to the owners about their items. They always ask questions like “Where did you get this item?” and “What can you tell me about its history?” Why do they ask these questions? Because the details in the answers to those questions provide knowledge about the authenticity and nature of the item. The dealer also carefully examines the item, looking for markings and dates that provide clues to the item’s origin.

Users of data must be able to put it into context before the data becomes useful as information. Information about data is referred to as *metadata*. The simplest definition of metadata is “data about data.” To be a bit more precise, metadata describes data, providing information such as type, length, textual description, and other characteristics. For example, metadata allows the user to know that the customer number is a five-digit numeric field, whereas the data itself might be 56789.

Metadata is “data about data.”

So, using our *Antiques Roadshow* example, the item being evaluated is the “data.” The answers to the antique dealer’s questions and the markings on the item are the “metadata.” Value is assigned to an item only after the metadata about that item is discovered and evaluated.

Metadata characterizes data. It is used to provide documentation such that data can be understood and more readily consumed by your organization.

Metadata answers the *who, what, when, where, why, and how* questions for users of the data.

From Data to Knowledge and Beyond

The basic building block of knowledge is data. *Data* is a fact represented as an item or event out of context and with no relation to other facts. Examples of data are 27, JAN, and 010110. Without additional details, we know nothing about any of these three pieces of data. Consider:

- Is 27 a number in base ten, or is it in octal (which would translate to 23 in base ten)?
- If 27 is a number in base ten, what does it represent? Is it an age, a dollar amount, an IQ, a shoe size, or something else entirely?
- What does JAN represent? Is it a woman's name (or a man's name)? Or does it represent the first month of the year? Or perhaps it is something else entirely?
- Finally, What about 010110? Is it a binary number? Or is it a representation of a date, perhaps January 1, 1910? January 1, 2010? Or something else entirely?

Data is a fact represented as an item or event out of context.

Because of the lack of context, these are all examples of data. *Information*, on the other hand, adds context by specifying relationships between data, and possibly other information. Data in context with metadata makes information. The relationships may represent information, yet the relations do not actually constitute information until they are understood. In addition, the relationships that represent data have a tendency to be limited in context, mostly about the past or present, with little if any implication for the future.

Webster's New Collegiate Dictionary defines *knowledge* as “the fact or condition of knowing something with familiarity gained through experience or association.” Knowledge adds understanding and retention to information. It is the next natural progression after information. To have “knowledge” requires information in conjunction with patterns found in data, information, and other knowledge. Therefore, knowledge couples information with understanding and cognition.

The final step would be to move from knowledge to wisdom. *Wisdom* can

be thought of as applied knowledge. You may have the knowledge that fatty foods are bad for you, but if you eat them anyway, you are not wise.

Wisdom can be thought of as applied knowledge.

In order for data to be anything more than simply data, metadata is required. Without metadata, data has no identifiable meaning—it is merely a collection of digits, characters, or bits. Metadata gives data its form and makes it usable by information professionals.

Metadata Strategy

A wise organization will develop a metadata strategy to collect, manage, and provide a vehicle for accessing metadata. A sound metadata strategy should address the following:

- A policy for how metadata is used in the organization
- Procedures for identifying and defining data ownership and stewardship
- Identification of the types of metadata that need to be collected
- A description of the purpose for each type of metadata that is identified—a clear and concise reason why each piece of metadata is required by the organization
- Methods for the collection and storage of metadata (typically using a repository)
- Methods for accessing the metadata
- Policies to enforce data stewardship procedures and security for metadata access
- Identification of metadata sources, both internal and external
- Measurements to gauge the quality and usability of metadata

Metadata publicizes and supports the data your organization produces and maintains. By assembling and managing metadata, your organization will have access to relevant facts about your data, making your systems more usable and your databases more useful.

DBAs should participate in the team that develops the metadata strategy, but the data administration organization, if one exists, should be the leader of the metadata effort.

Data Stewardship

A data steward is accountable for actions taken using a defined set of data. A data stewardship policy will identify individuals within the organization whose responsibility it is to create, maintain, and delete data. A data steward is not necessarily the data owner. A comprehensive data stewardship policy will also define the consumers of the data—that is, those who directly use the data during the course of their jobs.

Data Warehousing and Metadata

Companies that implement data warehousing systems are more likely than other companies to have embarked on a metadata management strategy. Users require accurate information about the data contained in a warehouse before the data can be used appropriately for business. Therefore, such businesses have a critical need for readily available, high-quality metadata. At times, though, little if any metadata is captured and managed prior to the onset of a data warehousing effort.

A data warehousing effort has a critical need for readily available, high-quality metadata.

Types of Metadata

Even though all metadata describes data, there are many different types and sources of metadata. At a fundamental level, though, all metadata is one of two types: technology metadata or business metadata. *Technology metadata* describes the technical aspects of the data as it relates to storing and managing the data in computerized systems. *Business metadata*, on the other hand, describes aspects of how the data is used by the business and is needed for the data to have value to the organization. Knowing, for instance, that the LICNO column is a positive integer between 1 and 9,999,999 is an example of technology metadata. Of course, the business user also requires this information. Knowing that a number referred to as a LICNO is the practitioner license number for certified course instructors, that it must be unique, and that every instructor can have one and only one license number is an example of business metadata (though these details also are also useful to the DBA in order to create the database appropriately and effectively).

All metadata is one of two types: technology metadata or business metadata.

For DBAs, the DBMS itself is a good source of metadata. The *system catalog* used to store information about database objects is a vital store of DBA metadata—technology metadata. DBAs and developers make regular use of the metadata in the DBMS system catalog to help them better understand database objects and the data contained therein. Depending on the DBMS, users can write queries against the system catalog tables or views, or they can execute system-provided stored procedures to return metadata from the system catalog tables. Just about any type of descriptive information about the composition of the data may be found in the system catalog. For example, most DBMSs store all of the following metadata in the system catalog:

- Names of every database, table, column, index, view, relationship, stored procedure, trigger, and so on
- Primary key for each table and any foreign keys that refer back to that primary key
- Which tables are in which views
- Data type, length, and constraints for each column of every table
- Names of the physical files used to store database data, as well as information about file storage, extents, and disk volumes
- Authorization and security information detailing which users have what type of authority on which database objects
- Date and time of the last database definition change, as well as the ID of the user who implemented the DDL for the change
- Database organization information

The DBMS system catalog is a particularly effective source of metadata because it is active, integrated, and nonsubvertible. The system catalog is *active* because the metadata is automatically built and maintained as database objects are created and modified. As the DBA creates databases, the DBMS automatically collects and populates metadata in the system catalog. The *integration* of the system catalog and the DBMS, coupled with the active nature of the system catalog, keeps the technology metadata in the system catalog accurate and up-to-date. Additionally, the DBMS system catalog is

nonsubvertible, meaning that normal DBMS operations are the only mechanism for populating the system catalog. Of course, the subvertibility of the system catalog will differ from DBMS to DBMS. Some DBMSs provide options to enable direct updates to the system catalog, but such an option is to be used only in emergencies and generally under the direction of the DBMS vendor's technical support personnel.

The DBMS system catalog is a particularly effective source of metadata.

Although a wealth of metadata can be found in the system catalog, this DBMS metadata is usually insufficient to fully describe data. For example, descriptions of database objects are not commonly found in the DBMS system catalog. Some DBMSs provide system catalog description columns that can be populated at the DBA's discretion. However, many DBAs avoid this for fear of disorganizing the system catalog. It's also possible that descriptions for the database objects were not available when the objects were created. Additional metadata that is useful but is not found in the system catalog includes

- Metadata for non-database files (flat or sequential files)
- Modification information regarding when and by whom data in the database was last changed
- Copybook information for the database table (or non-database file), as well as which programs use that information
- Information on batch jobs and transactions that access the data
- Operational metadata on IT infrastructure components
- Data model metadata describing the logical database design and how it maps to the physical database implementation
- Data warehousing and ETL metadata defining data source(s), system of record, and other analytical information
- Data ownership and stewardship metadata

Of course, this is an incomplete list. A myriad of different metadata types and purposes exist that can be cataloged and managed. Capturing and maintaining metadata better documents databases and systems, thereby making them easier to use. The more metadata that you make available to business users, the more value they will be able to extract from their

information systems.

Capturing and maintaining metadata makes databases and systems easier to use.

Repositories and Data Dictionaries

A [*repository*](#) stores information about an organization's data assets. In other words, repositories are used to store metadata. A properly implemented repository stores all pertinent metadata for the corporation. It can act as a single, centralized mechanism to assist in the migration of data from multiple sources to a data warehouse.

A [*repository*](#) stores all pertinent metadata for the corporation.

In choosing a repository, base your decision on the metadata storage and retrieval needs of your entire organization, not just the databases you wish to support. Typically, a repository can

- Store information about your data, processes, and environment.
- Support multiple ways of looking at the same data. An example of this concept is the three-schema approach, in which data is viewed at the conceptual, logical, and physical levels.
- Store in-depth documentation, and produce detail and management reports from that documentation.
- Support data model creation and administration. Integration with popular ETL, data modeling, and CASE tools is also an important evaluation criterion.
- Support for versioning and change control. Versioning helps to synchronize application development, eliminating rework and increasing flexibility.
- Enforce naming conventions.
- Parse and extract metadata from multiple sources. For example, if your site is a big COBOL shop, the repository vendor should offer tools that automatically examine your COBOL source code to extract metadata.
- Generate copybooks from data element definitions.

These are some of the more common functions of a repository. When

choosing a repository for database development, the following features generally are desirable:

- The data stores used by the repository can be stored using database tables in your DBMS. This enables your applications to directly read the data dictionary tables. For example, if you are primarily an Oracle shop, you should favor using a repository that stores its metadata information in Oracle tables. Some repository products use multiple DBMSs and allow the user to choose the DBMS to be used.
- The repository should be capable of directly reading the system catalog or views on the system catalog for each DBMS you use. This ensures that the repository will have current information on database objects.
- If the repository does not directly read the system catalog, an interface should be provided to simplify the task of populating the repository using the system catalog information.
- The repository provides an interface to any modeling and design tools used for the generation of database objects.

Most of the popular repository products are mainframe based and rely on a centralized metadata “database,” or repository. This approach is usually better suited for documenting OLTP-based systems. Such a repository may be more difficult to use in a data warehouse environment because a mainframe focus can present challenges when managing metadata in a distributed, state-of-the-art data warehouse implementation. Many ETL tools used in data warehousing projects also contain a repository that is geared toward the needs of the data warehouse. Organizations needing to manage metadata for both OLTP and data warehouses should make sure that the data in their ETL repositories can be migrated successfully to the OLTP repository.

Other repository products are application-centric. Such repository technology focuses on application development metadata—which is useful, but not comprehensive. For example, the Microsoft Repository is focused on Visual Studio and manages Microsoft computing assets.

Some repository products are application-centric.

Repository Benefits

Repository technology provides many benefits to organizations that properly

exploit their capabilities. The metadata in the repository can be used to integrate views of multiple systems, helping developers to understand how the data is used by those systems. Usage patterns can be analyzed to determine how data is related in ways that may not be formally understood within the organization. Discovery of such patterns can lead to business process innovation.

In general, the primary benefit of a repository is the consistency it provides in documenting data elements and business rules. The repository helps to unify the “islands of independent data” inherent in many legacy systems. The repository enables organizations to recognize the value in their legacy systems by documenting program and operational metadata that can be used to integrate the legacy systems with new application development.

A repository provides consistency in documenting data elements and business rules.

Furthermore, a repository can support a rapidly changing environment such as that imposed by Internet development efforts on organizations. The metadata in the repository can be examined to produce impact analysis reports to quickly determine how changes in one area will affect others.

Reusability is a big time-saver. If something can be reused instead of being developed again from scratch, not only will time be saved but also valuable resources can be deployed on more crucial projects. Repositories facilitate reuse by enabling the documentation of application components and making this metadata available to the organization.

Finally, repositories are an invaluable aid to data warehousing initiatives.

Repository Challenges

One of the biggest challenges in implementing and using repository technology is keeping the repository up-to-date. The repository must be populated using data from multiple sources—all of which can change at any time. When the composition or structure of source data changes, its metadata most likely will need to change, too.

Keeping the repository up-to-date can be difficult to accomplish.

The process for populating the repository is complicated and should be

automated as much as possible. Refer to [Figure 22.1](#). Metadata sources come from multiple areas and locations within an organization and can include

- Application component metadata from program development tools, application programs, and code libraries
- Business metadata from business user input, documents, and memos
- Data modeling metadata from data modeling tools
- Database metadata from the DBMS system catalog
- ETL metadata from data warehousing tools
- Operational metadata from automated operations and job scheduling tools
- Other types of metadata such as data usage metadata from query tools

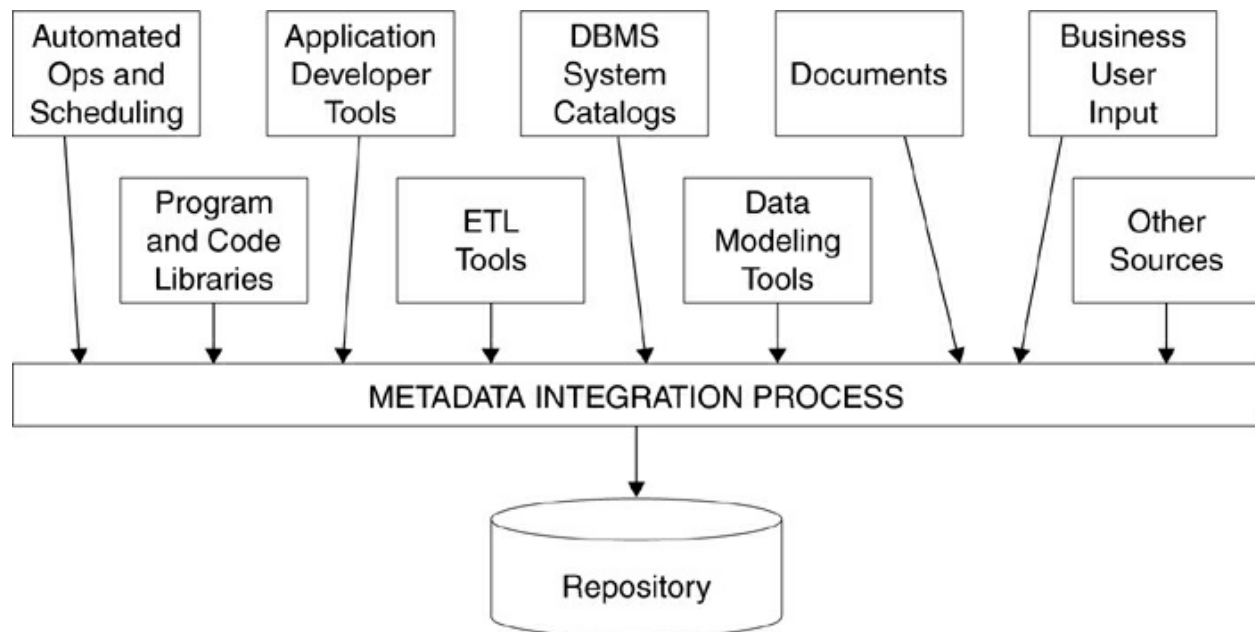


Figure 22.1. Populating the repository

Populating the repository is complicated.

To be successful, this information needs to be collected, parsed, and recorded in the corporate metadata repository. The integration process must take into account the frequency of change for each metadata source.

Whenever metadata changes at the source, the metadata in the repository will be out of sync until the source metadata is scanned, captured, and integrated into the repository again.

Many shops do not own a repository. More accurately, very few shops own a centralized metadata repository. Furthermore, many organizations that do own a repository do not always implement the proper integration and usage procedures, causing the repository to be neglected. As soon as the metadata in the repository becomes outdated, inaccurate, or nonexistent, the repository will cease to be of value. Of course, the fault does not necessarily lie with the repository technology—more likely the fault lies with the organization that does not implement procedures for keeping the metadata in the repository up-to-date. Of course, such an effort requires a significant budget, commitment, and the efforts of skilled data management professionals, including DAs and DBAs.

Very few shops own a centralized metadata repository.

Data Dictionaries

Data dictionaries were the precursors to repository technology. Data dictionaries were popular in the 1980s. The purpose of a [*data dictionary*](#) is to manage data definitions. In general, they offered little automation—the user had to manually key in the definitions. In some cases the data dictionary was integrated into the DBMS, and databases could be defined using the metadata in the data dictionary, but this was prerelational—before DBMS products had system catalogs.

The purpose of a [*data dictionary*](#) is to manage data definitions.

As more and more types of metadata were identified and organizations wished to accumulate and manage such metadata, the data dictionary was transformed into the repository. Use of CASE tools, such as Excelerator and Advantage Gen, for application and database development enabled more metadata to be captured and maintained during the development process. As developers became more sophisticated over time, data dictionaries evolved to provide more than just data attribute descriptions. The products became capable of tracking which applications accessed what databases. Developers who used the data dictionary properly were able to maintain their systems and applications more easily.

Truthfully, IBM's AD/Cycle and Repository Manager initiatives caused much of this transformation. Even though both initiatives ultimately failed in the marketplace, repository technology was forever changed by IBM's

ventures into this field. For more information on IBM's initiatives in this area, consult *IBM's Repository Manager/MVS* by Henry C. Lefkovits, the definitive book on the topic.

Summary

This chapter on metadata management has been necessarily brief. As a DBA, you will need to understand the role of metadata as it impacts the DBMS, databases, and database users. Organizations that spend a lot of time managing and maintaining metadata will likely have a data administrator on staff. Alternatively, the data warehouse administrator or architect might focus on metadata management.

DBAs may become involved in certain aspects of metadata management, such as repository selection, installation, and maintenance. However, most DBAs will *use* metadata far more than they will be called upon to store, manage, and maintain metadata.

DBAs will use metadata more than they will be called upon to store, manage, and maintain metadata.

Review

1. What is metadata?
2. Why is metadata important to an organization?
3. At a high level, what are the two types of metadata?
4. What is the difference between data and information?
5. Describe the purpose of a repository.
6. Why should the DBA participate in developing a metadata management strategy even in organizations with a well-defined data administration function?
7. Discuss the barriers to successful repository implementation.
8. Why is the DBMS system catalog useful for metadata analysis?
9. What is a data steward?
10. DBAs never require business metadata to do their job, and business users never need technology metadata: true or false? Explain why it is true or false.

Suggested Reading

- Baca, Murtha, ed. *Introduction to Metadata, Version 3*. Los Angeles, CA: Getty Publications (2008). ISBN 978-0-89236-896-9.
- Bruce, Thomas A. *Designing Quality Databases with IDEF1X Information Models*. New York, NY: Dorset House (1991). ISBN 0-932633-18-8.
- Carlis, John, and Joseph Maguire. *Mastering Data Modeling: A User-Driven Approach*. Boston, MA: Addison-Wesley (2001). ISBN 0-201-70045-X.
- Durell, William R. *Data Administration: A Practical Guide to Successful Data Management*. New York, NY: McGraw-Hill (1985). ISBN 0-07-018391-0.
- . *The Complete Guide to Data Modeling*. Princeton, NJ: Data Administration, Inc. (1993). No ISBN.
- Fleming, Candace, and Barbara von Halle. *Handbook of Relational Database Design*. Reading, MA: Addison-Wesley (1989). ISBN 0-201-11434-8.
- Hay, David C. *Data Model Patterns*. New York, NY: Dorset House (1996). ISBN 0-932633-29-3.
- Inmon, W. H. *Data Architecture: The Information Paradigm*. Wellesley, MA: QED Information Sciences (1989). ISBN 0-89435-268-7.
- Lefkovits, Henry C. *IBM's Repository Manager/MVS*. Wellesley, MA: QED Information Sciences (1991). ISBN 0-89435-349-7.
- Marco, David. *Building and Managing the Metadata Repository*. New York, NY: John Wiley & Sons (2000). ISBN 0-471-35523-2.
- Marco, David, and Michael Jennings. *Universal Meta Data Models*. New York, NY: John Wiley & Sons (2004). ISBN 0-471-08177-9.
- Modell, Martin E. *Data Analysis, Data Modeling, and Classification*. New York, NY: McGraw-Hill (1992). ISBN 0-07-042634-1.
- Perkinson, Richard C. *Data Analysis: The Key to Database Design*. Wellesley, MA: QED Information Sciences (1984). ISBN 0-89435-105-2.
- Ross, Ronald G. *Entity Modeling: Techniques and Application*. Boston, MA: Database Research Group (1988). ISBN 0-941049-00-0.
- Schmidt, Bob. *Data Modeling for Information Professionals*. Upper Saddle

- River, NJ: Prentice Hall (1999). ISBN 0-13-080450-9.
- Seiner, Robert S. "Questions Metadata Can Answer." *The Data Administration Newsletter*. <http://www.tdan.com/i011fe01.htm>.
- Simsion, Graeme C., and Graham C. Witt. *Data Modeling Essentials*. 3rd ed. San Francisco, CA: Morgan Kaufmann (2005). ISBN 0-12-644551-6.
- Tannenbaum, Adrienne. *Implementing a Corporate Repository: The Model Meets Reality*. New York, NY: Wiley Professional Computing (1994). ISBN 0-471-58537-8.
- . *Metadata Solutions*. Boston, MA: Addison-Wesley (2002). ISBN 0-201-71976-2.
- Wertz, Charles K. *The Data Dictionary: Concepts and Uses*. Wellesley, MA: QED Information Sciences (1986). ISBN 0-89435-180-X.

23. DBA Tools

Every major DBMS product provides a complete, functional database management system that can be used out of the box to store and manage data. Although organizations can install and use the DBMS as delivered, many will quickly find that the functionality needed to adequately support large-scale database development is not provided by the DBMS product alone.

The administration and maintenance of database applications is time-consuming if you use only the standard features of the DBMS. Fortunately, many DBA tools that enhance the functionality of relational database management systems are available from third-party vendors.

Types and Benefits of DBA Tools

A DBA tool reduces the amount of time, effort, and human error involved in maintaining efficient database systems and applications. Such tools ease the administrative burden and reduce the possibility of error.

DBA tools reduce the amount of time, effort, and human error involved in maintaining efficient database systems.

The need for these tools becomes apparent simply by reviewing the sheer number of products that are available. Most organizations implement at least one add-on tool for their DBMS; many organizations use dozens to ease the administrative burdens associated with managing production database development. Many varieties of tools are available that fulfill market niches not adequately supported by the major DBMS vendors. The remainder of this chapter provides a rundown of the categories and types of products available to the DBA to assist in managing and administering databases.

Data Modeling and Design

Database modeling and design tools provide a consistent and coherent means of creating conceptual and logical data models and transforming them into physical database designs. Database modeling and design tools do not have to be unique to a specific database, but some are. Tools developed specifically to support your DBMS can significantly reduce development time by automating repetitive tasks and validating the models. However, if you use

multiple DBMS products in your organization, you will be better off choosing a tool that can support all of them rather than choosing multiple tools geared for a specific database.

Tools specifically supporting your DBMS can significantly reduce development time.

When choosing a modeling and design tool, look for one that can

- Support the standard tasks associated with logical data modeling such as entity-relationship diagramming and normalization.
- Create a physical data model geared to each of your target DBMS platforms. This model should support all features of each DBMS. For example, for DB2 on the mainframe it should be able to depict all DB2 objects, referential integrity, VCAT (volume catalog)- and STOGROUP-defined tablespaces, system-managed storage, and capacity planning.
- Provide an expert system to verify the accuracy of the physical data model and to suggest alternative solutions.
- Cross-reference the logical model to the physical model, capturing text that supports physical design decisions such as denormalization and type of tablespace.
- Generate standard DDL automatically to fully implement the database defined in the physical data model. The DDL created by the tool from the physical model should support all of the options for each DBMS supported.
- Interface with application development tools and repository products used within the organization.

Database Change Management

As we discussed in [Chapter 7](#), making changes to your databases can be an arduous and error-prone job. However, it is a rare database that does not need to undergo some sort of change during its useful life. A variety of tools are available to assist the DBA in managing and effecting database changes. These tools optimize and automate multiple change management tasks, including database alteration, database comparison, security authorization, audit tracking, catalog querying, space management, and handling referential

integrity.

A variety of tools are available to assist in managing and effecting database changes.

Change Management Tools

The predominant form of change management tool is the database alteration and comparison tool. Although the structure of relational databases can be modified using the ALTER statement, this statement is functionally crippled in most DBMS products. In theory, the DBA should be able to alter all of the parameters that can be specified for an object when it is created, but no current DBMS product supports this. For example, most DBMS products enable you to add columns to an existing table, but only at the end. Furthermore, they do not permit the DBA to remove columns from a table—instead, the table must be dropped and then recreated without the specified columns.

The ALTER statement is functionally crippled in most DBMS products.

Another problem that DBAs encounter when modifying database structures is the cascading DROP effect. If a change to a database object mandates that it be dropped and recreated, all dependent objects are dropped when the database object is dropped. This includes tables, all indexes on the tables, all primary and foreign keys, any related synonyms and views, any triggers, all authorizations, and, of course, the data. Ensuring that DDL is issued after the modification to reverse the effects of cascading DROPs can be a tedious, complex, and error-prone procedure.

Many types of database object alteration cannot be performed using the generic ALTER statement.

Additionally, many other types of database object alterations cannot be performed using the generic ALTER statement. Depending on the DBMS, you may not be able to

- Change the name of a database, table, alias, view, column, tablespace, dbspace, trigger, stored procedure, user-defined function, relationship, or index
- Change database partitioning

- Move a table from one database, dbspace, or tablespace to another
- Rearrange column ordering
- Change a column's data type and length
- Remove columns from a table
- Change the primary key without dropping and adding the primary key
- Add columns to a view or remove columns from a view
- Change the SELECT statement on which a view is based
- Change the indexing columns
- Change the uniqueness specification of an index
- Change the clustering of an index or table

Such a list usually provides all the justification needed to obtain a database alteration tool. Of course, the exact list will differ from DBMS to DBMS. Database alteration tools provide an integrated environment for altering database objects. Such tools usually provide a menu-driven or point-and-click interface that allows the DBA to specify the type of change needed. The burden of ensuring that the database change is performed correctly is moved from the DBA to the tool. At a minimum, an ALTER tool should

The burden of ensuring that the database change is performed correctly is moved from the DBA to the tool.

- Maintain tables easily without manually requiring the DBA to code DDL.
- Retain or reapply all dependent objects, authorizations, and data affected by the ALTER if a drop is required.
- Navigate hierarchically from object to object.
- Provide GUI-based modification showing “before” and “after” definitions of the objects before the changes are applied.
- Batch requested changes into a group that can be executed in the foreground or the background.
- Analyze changes to ensure that the requested alterations do not violate any DDL rules. For example, if a series of changes is requested and one change causes a subsequent change to be invalid (an object is dropped, for instance), this should be flagged before execution.

- Provide the capability to monitor changes as they are applied.

An automated change management solution enables the DBA to focus on the required change instead of on the arcane details of how the DBMS implements such change. The tool is built to understand not just the discipline of change management but also the DBMS in which the changes are to be made. This built-in intelligence shifts the burden of ensuring that a change to a database object does not cause other implicit changes from the DBA to the tool. Furthermore, once the change has been identified and implemented for one system, it can easily be deployed on other database copies with minimal, or perhaps no, changes.

Another benefit of a change management tool is in database analysis and planning. The impact of changes can be examined prior to implementing any change. This is an invaluable resource for ensuring safe and efficient database changes. This type of tool also uses automation to minimize the resources required to implement database change. Instead of writing a new, complex change script from scratch for each database change, the DBA can rely on the change management tool to accomplish this. Application and database availability will be enhanced because the product will implement the change in the least intrusive, quickest manner possible.

The impact of changes can be examined prior to implementing any change.

All in all, a database change management product will improve availability, minimize errors, and speed up your time to market.

Database Comparison Tools

Over time, the DBA will make changes to databases throughout the organization. It is possible that even with an automated database change management solution, some changes will be implemented on some systems but not on others. DBAs at large companies must keep track of dozens, if not hundreds or thousands, of database servers. Ensuring that changes are migrated effectively to all of those database servers can be a difficult task.

Ensuring that changes are migrated effectively can be a difficult task.

A database comparison tool enables the DBA to compare one database to

another in terms of its database objects and structures. Such tools will identify differences and automatically generate the DDL to make the databases the same—from a structural perspective, not a data content perspective. The database comparison tool should allow the DBA to make the following types of comparisons:

- One live database to another live database (on the same server or another server)
- A live database to a DDL script file
- One DDL script file to another DDL script file

DBAs who manage a large number of database servers should consider using such products. As mentioned earlier, database comparison tools compare only the structure of the database, not the contents. However, some vendors offer tools that can compare the contents of one table to the contents of another. Such tools are often useful during application program testing and debugging.

Database Object Migration Tools

Many DBMS products do not provide a feature to migrate database objects from one database server or subsystem to another. Without a tool, you can accomplish this only by manually storing the DDL CREATE statements (and all subsequent ALTER statements) in a script file and then executing the script on the other database server. The DBA would then need to unload the data from the source system and load it to the target system (if the data were being migrated as well as the database schema). Manual processes such as this are error prone. In addition, this process does not take into account the migration of table data and database security.

Migration tools facilitate the quick migration of database objects from one environment to another (e.g., from test to production). Although similar to a table altering tool, an object migration tool has minimal altering capability (some interface directly with an ALTER tool or are integrated into a single tool). The migration procedure is typically driven by panels or GUI screens that prompt the user for the objects to migrate.

Migration tools facilitate the quick migration of database objects.

Migration can typically be specified at any level. For example, if you

request the migration of a specific database, you also could migrate all dependent objects and security. Capability is provided such that database object names, authorization IDs, and other objects can be renamed according to the standards of the receiving instance, subsystem, or server. When the parameters of the migration have been specified completely, the tool creates a job to implement the requested database objects in the requested environment.

A migration tool can reduce the time required by database administrators to move databases from environment to environment. Quicker turnaround results in a more rapid response to user needs, thereby increasing the efficiency of your business.

Referential Integrity Tools

Referential integrity (RI) is used by relational databases to ensure the validity of primary key to foreign key relationships. However, RI can be difficult to administer and implement. RI tools eliminate this difficulty by

- Analyzing data for both system- and user-managed referential integrity constraint violations
- Executing faster than the DBMS-provided integrity checking facility or utility
- Enabling additional types of RI to be supported; for example, by analyzing primary keys for which no foreign keys exist and deleting the primary key row (pendant DELETE processing)

Additionally, tools are available that allow data to be viewed and extracted in referential sets. Such a capability makes it easier to create reasonable test data by using a subset of the data in the production databases.

Catalog Query and Analysis Tools

The system catalog or data dictionary contains a wealth of information essential to the operation of the DBMS. Information about all database objects, authority, and recovery is stored and maintained in the system catalog. DBAs rely on this information to perform their job. The system catalog is composed of relational tables and can be queried using SQL and/or vendor-supplied stored procedures. Some DBMS vendors provide system catalog views that are easier to query and monitor. However they are accessed, these tables provide an information base for many monitoring and

administrative tasks.

DBAs rely on the system catalog to perform their job.

Coding SQL every time the DBA needs to access information in the system catalog can be a very time-consuming process. Often, the DBA must combine information from multiple catalog tables to provide the user with facts relevant for a particular task. Moreover, for the most part, DBMS vendors have made the catalog tables difficult to understand and query by using odd naming conventions, denormalized structures, unused columns, poor data type choices, and little documentation. As DBMS vendors add features to new versions of their products, the system catalog becomes more and more difficult to understand as new data is forced into an already ugly system catalog design and implementation. When stored procedures and views are provided, querying the catalog tables is easier; however, crucial information is sometimes missing from these canned “queries.”

Add-on query tools can ease the burden of developing SQL queries to access the system catalog tables. Sometimes these tools are referred to as *catalog visibility tools* because they make it easier to access information stored in the system catalog. The basic feature common to all catalog tools is the capability to request catalog information using a GUI (or panel-driven) interface without using SQL statements.

A GUI is a basic feature common to all catalog tools.

System catalog tools that provide only this level of capability are rudimentary tools at best. Most of these tools provide much more functionality. Instead of merely enabling data access, many tools can do one or more of the following tasks:

- Create syntactically correct DDL statements for all database objects by reading the appropriate system catalog tables. These statements are generally executed immediately or saved in a data set for future reference or use.
- Modify any updateable statistical columns using a non-SQL interface.
- Create syntactically correct authorization/security statements from the catalog in the same way that DDL is generated.
- Perform “drop analysis” on a SQL DROP statement. This analysis

determines the effect of the cascading DROP by detailing all dependent objects and security that will be deleted as a result of executing the DROP.

- Provide a hierarchic listing of database objects. For example, if a specific table is chosen, the tool can migrate quickly up the hierarchy to show its tablespace and database, or down the hierarchy to show all dependent indexes, views, synonyms, aliases, referentially connected tables, and plans.
- CREATE and DROP database objects, and grant and revoke security from a screen without coding SQL. Additionally, some tools log all DROPS and REVOKEs such that they can be undone in the event of an inadvertent DROP or REVOKE execution.
- Operate directly on the system catalog or on a copy of the system catalog to reduce system-wide contention.
- Generate the full set of database commands used to manage the database. Examples could include starting/stopping a database or tablespace, displaying system information, and many other commands.

These features aid DBAs in performing their day-to-day duties.

Furthermore, a catalog query tool can greatly diminish the amount of time required for a junior DBA (or new staff member) to become a productive member of the DBA team.

Table Editors

There are only two methods of updating relational data supplied by most DBMS products out of the box:

- SQL DELETE, INSERT, and UPDATE statements
- Database utilities such as LOAD or IMPORT

SQL statements operate on data a set at a time, so a single SQL statement can affect multiple rows—or even all of the rows. Coding SQL statements for every data modification required during the application development and testing phase can be time-consuming. Furthermore, database utilities such as LOAD and IMPORT are not a viable means of making small targeted changes to data. They are designed and optimized for moving data in bulk.

A table editing tool can reduce the time needed to make simple data alterations by providing full-screen editing capability for database tables. The

user specifies the table to edit, and a table editor launches. The data is presented to the user as a series of rows, with the columns separated by spaces. A header line indicates the column names. The data can be scrolled up and down as well as left and right. To change data, the user simply types over the current data.

A table editing tool can reduce the time needed to make simple data alterations.

This type of tool is ideal for supporting the application development process. A programmer can make quick changes without coding SQL. In addition, if properly implemented, a table editor can reduce the number of erroneous data modifications made by beginning SQL users.

Exercise caution before using a table editor on critical production data. When a table editor is used, all columns are available for update, and a simple miskeying can cause unwanted updates. Native SQL should be used if you must ensure that only certain columns are updated. Tested SQL statements and application programs are characterized by their planned nature—the modification requests were well thought out and tested. This is not true for changes implemented through a table editor.

Performance Management

Assuring optimal performance is one of the biggest problems faced by DBAs on an ongoing basis. The loudest complaints come from users who have to wait longer than they are used to waiting for their applications to respond. This is especially true if the users never had to wait in the past.

However, what causes those formerly fast applications to stall and deliver subpar performance? If there were an easy answer to that question, many DBAs would be out of work. Performance management tools help the DBA to gauge the responsiveness and efficiency of SQL queries, database structures, and system parameters. Such tools run in the background to capture database performance statistics and alert the DBA when problems occur. Advanced performance tools can take proactive measures to correct problems as they happen.

Every database application, at its core, requires three components in order to operate: the system, the database, and the application. To deliver performance, the DBA must be able to monitor and tune each of these

components. There are tools available to the DBA for monitoring and optimizing each of these components.

To deliver performance, the DBA must be able to monitor the system, the database, and the application.

System Performance Tools

System performance tools examine the database server, its configuration, and usage. The most commonly used system performance tool is the *performance monitor*. Database performance monitoring and analysis tools support many types of performance-oriented requests in many ways. For example, system performance tools can operate

- In background mode as a batch job that reports on performance statistics written by the DBMS trace facility
- In foreground mode as an online monitor that either traps trace information or captures information from the DBMS control blocks as applications execute
- By sampling the database kernel and user address spaces as the program runs and by capturing information about the performance of the job, independent of database traces
- By capturing database trace information and maintaining it in a history file (or table) for producing historical performance reports and for predicting performance trends
- As a capacity planning device that gives statistical information about an application and the environment in which it will operate
- As an after-the-fact analysis tool on a workstation that analyzes and graphs all aspects of application performance and system-wide performance

Each database performance monitor supports one or more of these features. The evaluation of database performance monitors is a complex task. Sometimes more than one performance monitor is used at a single site—perhaps one for batch reporting and another for online event monitoring. Maybe an enterprise-wide monitoring solution has been implemented and one component of that solution is a database module that monitors your DBMS, but it lacks the details of a more sophisticated DBMS monitor. So, another performance monitor is purchased for daily DBA usage, while the module of

the enterprise-wide monitoring solution is used for integrated monitoring by system administrators.

Modern database performance tools can set performance thresholds that, once reached, will alert the DBA, perform another task to report on, or actually fix the problem. These tools are typically agent based. An *agent* is a piece of independent code that runs on the database server looking for problems. It interacts with, but does not rely on, a console running on another machine that is viewed by the DBA. This agent architecture enables efficient database monitoring because the agent is not tied to a workstation and can act independently. The agent sends information to the DBA only when required.

Additionally, some system performance tools are available that focus on a specific component of the DBMS such as the data cache. Such a tool can be used to model the memory requirements for database caching, to capture data cache utilization statistics, and perhaps even to make recommendations for improving the performance of the data cache.

Some system performance tools focus on a specific component of the DBMS such as the data cache.

Another type of performance optimization tool enables database configuration parameters to be changed without recycling the DBMS instance, subsystem, or server. These tools are useful when the changes require the DBMS to be stopped and restarted. Such tools can dramatically improve availability, especially if configuration parameters need to be changed frequently and the DBMS does not support dynamic parameter modification.

A few ISVs provide *invasive system performance tools* that enhance the performance of databases by adding functionality directly to the DBMS and interacting with the database kernel. Typically, these products take advantage of known DBMS shortcomings.

For example, products are available that enhance the performance of reading a database page or block or that optimize data caching by providing additional storage and control over buffers and their processing. Care must be taken when evaluating invasive performance tools. New releases of the DBMS may negate the need for these tools because functionality has been added or known shortcomings have been corrected. However, this does not mean that you should not consider invasive database performance tools. They

can pay for themselves after only a short period of time. Discarding the tool when the DBMS supports its functionality is not a problem if the tool has already paid for itself in terms of better performance.

Care must be taken when evaluating invasive performance tools.

One final caution: Because invasive performance tools can interact very closely with the database kernel, be careful when migrating to a new DBMS release or a new release of the tool. Extra testing should be performed with these tools because of their intrusive nature.

Database Performance Tools

Most DBMSs do not provide an intelligent database analysis capability. Instead, the DBA or performance analyst must use system catalog views and queries, or a system catalog tool, to keep watch over each database and its objects. This is not an optimal solution because it relies on human intervention for efficient database organization, opening up the possibility for human error.

Fortunately, database analysis tools are available that can proactively and automatically monitor your database environment. These database analysis tools typically can

- Collect statistics for tables and indexes: standard statistical information, extended statistics capturing more information (for example, data set extents), or a combination of both
- Read the underlying data sets for the database objects to capture current statistics, read the database statistics from the system catalog, read tables unique to the tool that captured the enhanced statistics, or any combination thereof
- Set thresholds based on database statistics whereby the automatic scheduling of database reorganization and other maintenance tasks can be invoked
- Provide a series of canned reports detailing the potential problems for specific database objects

Application Performance Tools

Writing SQL statements to access database tables is the responsibility of an

application development team. However, the DBA usually gets involved when it comes to the performance of SQL. With SQL's flexibility, the same request can be made in different ways. Because many of these methods are inefficient, application performance can fluctuate wildly unless the SQL is analyzed and tuned by an expert prior to implementation.

On the Importance of Application Performance . . .

As much as 80 percent of all database performance problems are caused by inefficient SQL and application code.

The EXPLAIN or SHOWPLAN commands provide information about the access paths used by SQL queries by parsing SQL in application programs and placing encoded output into a PLAN_TABLE or by producing a standard access path report. To gauge efficiency, a DBA must decode this data and determine whether a more efficient access path is available.

SQL code reviews are required to ensure that optimal SQL design techniques are used. As discussed in [Chapter 6](#), an application walk-through should perform a review of all SQL statements, the selected access paths, and the program code in which the SQL is embedded. The review also includes an evaluation of database statistical information to ascertain whether production-level statistics were used at the time of the EXPLAIN.

A line-by-line review of application source code and EXPLAIN output is tedious and prone to error, and it can cause application backlogs. SQL analysis tools greatly simplify this process by automating major portions of the code review process. The SQL analysis tool typically

- Analyzes the SQL in an application program, describing the access paths chosen in a graphic format, an English description, or both.
- Issues warnings when specific SQL constructs are encountered. For example, each time a sort is requested (by ORDER BY, GROUP BY, or DISTINCT), a message informs the user of the requisite sort.
- Suggests alternative SQL solutions based on an “expert system” that reads SQL statements and their corresponding PLAN_TABLE entries and poses alternative SQL options.
- Extends the rules used by the “expert system” to capture site-specific rules.

- Analyzes at the subsystem, instance, server, application, plan, package, or SQL statement level.
- Stores multiple versions of EXPLAIN output, creates performance comparisons, and plans history reports.

SQL analysis tools can automate major portions of the code review process.

Tools that analyze the performance of the application code in which the SQL is embedded are available too. These tools usually capture in-depth information about programs as they are run and provide reports that specify which areas of the code consume the most resources. Unfortunately, most of these tools do not necessarily interface to SQL analysis tools. If you are interested in why this might be a problem, see the sidebar “Application and SQL Analysis.”

Application and SQL Analysis

Consider an application program that contains a singleton SELECT inside a loop. The singleton SELECT requests a single row based on a WHERE clause, checking for the primary key of that table. For each iteration of the loop, the program changes the primary key value being searched such that the entire table is read from the lowest key value to the highest key value.

SQL analysis tools will probably not target the SQL statement as inefficient because the predicate value is for the primary key, which should invoke indexed access. The application program analysis tool may flag the section of the code that accesses the data as inefficient, but it will not help you to fix it or tell you why it is inefficient.

A knowledgeable performance analyst or DBA would have to use both tools and interpret the output of each to arrive at a satisfactory conclusion. For example, it could be more efficient to code a cursor, without a predicate, to retrieve every row of the table, and then fetch each row one by one. This method would eliminate index I/O, might use parallel access, and therefore should reduce I/O and elapsed time—thereby enhancing performance.

Only a trained analyst can catch this type of design problem during a code walk-through. Although a plan analysis tool significantly reduces

the effort involved in the code review process, it cannot eliminate it.

The first feature required of SQL analysis tools is the ability to read and interpret standard EXPLAIN or SHOWPLAN output. The tool should be able to read the plan table or interface directly with the DBMS command to obtain the output. It then must be able to automatically scan the EXPLAIN or SHOWPLAN data and report on the selected access paths and the predicted performance. Advanced tools will provide recommendations for improving the SQL by adding indexes or modifying the SQL.

End-to-End Performance Tools

Modern applications require multiple system components and run across multiple networked devices to deliver functionality. When performance problems arise, it can be difficult to determine what, exactly, is causing the problem. Is it on the client or the server? Is it a networking problem? Is it a database issue or a code problem?

End-to-end performance monitoring tools exist that track an application request from initiation to completion. These solutions provide enhanced visibility specifically into application performance—giving organizations the power to understand both when and why performance has degraded, and the information needed to improve matters in a business-prioritized fashion.

End-to-end performance monitoring tools track application requests from initiation to completion.

Following the workload as it progresses across multiple pieces of hardware and software makes problem determination possible.

Backup and Recovery

As we learned in [Chapters 15](#) and [16](#), assuring the recoverability of database systems is a complex task. In order to do the job properly, DBAs need to understand the backup and recovery features of the DBMS, how it interfaces to the disk storage systems, and the business impact of data to the organization. Sometimes DBAs need a helping hand.

Fortunately, there are a number of tools that streamline and automate the backup and recovery process. The simplest form of backup and recovery tool is the high-speed utility. ISV utilities that speed up the process of making

image copy backups and using those backups for recovery can be used to reduce downtime and increase data availability. Some tools also simplify, not just speed up, these processes. For example, many high-speed backup and recovery utilities can back up entire databases or database objects using masking and wildcarding techniques. For example, consider the following command:

```
COPY GL21DBX2.T*
```

A command such as this might be used to make an image copy backup of every tablespace in the database named GL21DBX2 that begins with the letter *T*. Not only is this simpler than explicitly listing every tablespace, but the backup script will not need to be edited if a new tablespace is added that conforms to the wildcarded mask.

The simplest form of backup and recovery tool is the high-speed utility.

Additional products are available that automate the entire recovery process. Such tools might be able to accept a recovery request, examine the log and the image copy backups, make a recommendation on how to recover, or, even further, build the scripts to perform the recovery that would result in the most data recovered in the least amount of downtime.

A final type of recovery tool is the *log-based recovery tool*. Such a tool can be used to examine the log and produce reverse SQL, as discussed in [Chapter 15](#). For example, assume the DBA issued a DELETE statement in error. Using the log-based recovery tool, the DBA would input the program name and the time it ran. The tool would examine the log, find the DELETE, and create INSERT statements to reproduce the deleted data. Simply by running the INSERT statements, the database is recovered. Of course, this is a simplistic example, but a log-based analysis tool can come in handy when erroneous database modifications need to be identified and rectified quickly.

Database Utilities

Many database utilities that ship free with the DBMS are simple, no-frills programs that are notorious for poor performance, especially on very large tables. However, these utilities are required to populate, administer, and organize your databases. The typical utilities that are provided are LOAD, UNLOAD, REORG, BACKUP, and RECOVER, as well as utilities for

integrity checking.

Many database utilities that ship free with the DBMS are notorious for poor performance.

Third-party vendors provide support tools that replace the database utilities and provide the same or more functionality in a more efficient manner. For example, it is not unheard of for third-party vendors to claim that their utilities execute anywhere from four to ten times faster than the native DBMS utilities. These claims must be substantiated for the data and applications at your organization (but such claims are quite believable). Before committing to any third-party utility, the DBA should be sure that the product conforms to at least the following requirements:

- Does not subvert the integrity of the data in the database.
- Provides at a minimum the same features as the corresponding native utility. For example, if the REORG utility reorganizes both indexes and tablespaces, the enhanced REORG tool must be able to do the same.
- Does not subvert standard database features. For example, DB2 image copies are maintained in the DB2 system catalog. An enhanced backup tool for DB2, therefore, should store its image copies there as well.
- Provides an execution time at least twice as fast as the corresponding database utility. For example, if the Sybase DUMP utility requires 20 minutes to back up a table, the enhanced backup tool must dump the same table in no more than 10 minutes. Of course, this does not need to be a hard-and-fast rule. Sometimes even a moderate increase in processing time is sufficient to justify the cost of a third-party utility offering.
- Corrects deficiencies in the standard database utilities. For example, if the LOAD utility does not load data in sequence by clustering index, the replacement LOAD utility should do so.

When testing utility tools from different vendors, be sure to conduct fair tests. For example, always reload or recover prior to testing REORG utilities, otherwise you may skew your results due to different levels of table organization. Additionally, always run the tests for each tool on the same object with the same amount of data, and make sure that the data cache is flushed between each test run. Finally, make sure that the workload on the

system is the same (or as close as possible) when testing each product because concurrent workload can skew benchmark test results.

Conduct fair tests when testing utility tools from different vendors.

Utility Management Tools

Another variety of database utility offering is the utility manager. This type of tool provides administrative support for the creation and execution of database utility job streams. These utility generation and management tools

- Automatically generate utility parameters, JCL, or command scripts.
- Monitor the database utilities as they execute.
- Automatically schedule utilities when exceptions are triggered.
- Restart utilities with a minimum of intervention. For example, if a utility cannot be restarted, the utility manager should automatically terminate the utility before resubmitting it.

Data Protection, Governance, Risk, and Compliance Tools

As we discussed in [Chapters 15](#) and [16](#), database security and protection, along with governance, risk, and compliance (GRC) requirements, impose the need for security and compliance management. There are tools that can be used to improve database security and compliance, some of which are important for DBAs to understand and use.

On the other hand, GRC is a business proposition, and the majority of the solutions in this category are not relevant for DBAs. The DBA must be aware of the tools that integrate with the DBMS and their databases, or impose additional requirements in terms of administrative workload or performance implications.

Auditing Tools

An *audit* is the examination of a practice to determine its correctness. Database auditing software helps in monitoring the data control, data definition, and data integrity in the database environment. Most DBMS products provide limited audit mechanisms, but these features are typically difficult to use and maintain. Auditing tools provide the ability to audit at a more granular level and simplify the ability to report on the audit trails.

Auditing tools provide the ability to audit at a granular level.

A typical auditing facility permits auditing at different levels within the DBMS, for example, at the database, database object level, and user levels. But capturing so much information, particularly in a busy system, can cause performance to suffer. Production of the required audit details must be accomplished without diminishing the operations of the computerized systems that keep the practice functioning.

The detail and authenticity of the audit trail produced are just as important as the operational systems' performance. The audit trails must be detailed enough to capture "before" and "after" images of database changes. If the mechanism capturing the audit details is not comprehensive and efficiently engineered, it ceases to be a compliance solution. Furthermore, the audit trails must be stored somewhere that protects the authenticity of the audited information while allowing seamless access for reporting.

Due to the potential volume of changes made to database data, a useful auditing facility must allow for the selective creation of audit records to minimize performance and storage problems. The general rule of thumb is that only data that must be audited to be in compliance should be audited, and nothing more.

Typical aspects that can be captured by database auditing tools include

- All accesses by high-level users (such as DBADMIN and SYSADMIN)
- Authorization failures
- SQL GRANT and REVOKE statements
- DDL statements issued
- DML statements issued
- SQL BIND requests or new program invocations
- Utility executions

There are several popular techniques that can be deployed to audit database data. By far, the best technique engages proactive monitoring of database operations directly at the database server. This technique captures *all* requests for data as they are made. By capturing the audit details at the server level, the software can guarantee that all access is monitored. Other techniques, such as trace-based auditing or parsing database logs, can miss certain types of database activities.

A robust database access auditing solution that addresses regulatory compliance should be able to provide answers to at least the following questions:

1. Who accessed the data?
2. At what date and time was the access?
3. What program or client software was used to access the data?
4. From what location was the request issued?
5. What SQL was issued to access the data?
6. Was the request successful, and if so, how many rows of data were retrieved?
7. If the request was a modification, what data was changed? (A “before” and “after” image of the change should be accessible.)

Of course, there are numerous details behind each of these questions. A robust database access auditing solution should provide an independent mechanism for long-term storage and access of audit details. The solution should offer canned queries for the most common types of queries, but the audit information should be accessible using industry-standard query tools to make it easier for auditors to customize queries as necessary.

The auditing tool should produce a set of prepackaged reports geared for supporting industry and governmental regulations. For example, to support PCI DSS, the tool should offer prepackaged reports geared to prove compliance. The reports should be easy to produce and easy to read.

Organizations with strict security and auditing requirements, or with significant regulatory or industry compliance requirements, should consider using a database auditing tool because of the weak auditing capabilities of most DBMS products.

Access Trending Tools

Trending tools usually rely on database auditing technology to watch database accesses and establish trends and patterns. Such tools are useful for identifying and prohibiting behavior outside of the norm to proactively halt database attacks.

As an example, consider an end user who works on sensitive data on a daily basis. A trending tool can establish that user’s normal patterns. Perhaps the person generally accesses sensitive data between 9:00 and 11:00 in the

morning and again between 2:00 and 4:00 in the afternoon. The data is accessed methodically, usually on a customer-by-customer basis. And the user does this consistently, more or less, every workday.

Then, on Saturday evening, the same user attempts to access sensitive data for all of the customers in the Northeast region. That sounds like a potential security issue. The trending tool can capture and report upon this so that the atypical request is documented. Or, for organizations that are more cautious, the trending tool may be able to prohibit such nonstandard access, requiring management authorization before allowing it to proceed.

Data access trending has the potential to protect data that you may not even know is being surreptitiously accessed.

Data Profiling Tools

Another useful category of data management tool provides data profiling capabilities. Profiling your data is a methodology for gaining insight into your business data and refining your processes for improving data quality.

A data profiling tool is used to discover structure, data, and relationship anomalies in your data. Profiling techniques help you to uncover instances of improper data where the data does not match the metadata definition, patterns do not match, values are inaccurate, and there are occurrences of redundant data.

For example, consider a database table with a column containing Social Security numbers. We know that they should all be of the format nnn-nn-nnnn. A data profiling tool can be deployed to examine that particular column to show what percentage of the data does not conform to that pattern. Or, consider another column where you do not know the specific pattern of the data. In this case, the data profiling tool can uncover a pattern where perhaps one was thought not to exist.

At any rate, data profiling tools are useful for cleaning up data, such as is required for compliance and data warehousing projects.

Data Masking Tools

Data masking is an additional category of data protection solution that can be used to obfuscate actual values in the database. You may need to mask data values for reporting or other reasons to hide sensitive data. PCI DSS regulations are one area that requires data masking. Payment card values

cannot appear unmasked on receipts. The next time you use a credit card to buy something, try to find your credit card number. It will be masked.

An additional place where data masking may be required is when accessing production data for testing purposes. The testing process requires valid data, but it does not need to be actual values. Data masking solutions can process “real” values and turn them into valid but different values for testing.

For example, you might need to mask Social Security numbers. You do not want personally identifiable details such as name, address, and so on to be associated with an actual Social Security number on most reports.

Security Tools

Database security is usually provided internal to the database by use of GRANT and REVOKE SQL statements, which grant authorization explicitly and implicitly to database users. Some DBMS products provide authorization exits to enable communication with external security management packages. This eases the administrative burden of database security by enabling the corporate data security function to administer groups of users.

The implementation of security in most relational databases has several problems. Paramount among these deficiencies is the effect of the cascading REVOKE. If authority is revoked from one user who has previously granted authority to other users, all dependent authorizations also are revoked.

A database security add-on tool can address this problem. These tools typically analyze the effects of a REVOKE. These tools enable the user to revoke the authority and optionally reassign all dependent authority, either by recreating the appropriate GRANT statements to reapply the authorizations implicitly revoked or by revoking the authority and automatically reapplying all implicit REVOKEs in the background.

A database security add-on tool typically analyzes the effects of a REVOKE.

Database security tools provide other functions. Consider the administrative overhead when database users are hired, quit, or are transferred. Security must be added or removed. A good security tool enables a user to issue a GRANT LIKE command that can copy database authority from one database object to another or from one database user to another.

Suppose that a DBA is transferred to another department. A security tool can assign all of that DBA's authority to another user before revoking his authority. Or suppose that a new table is created for an existing application, and it requires the same authorizations as the other tables in the application. This type of tool enables a user to copy all security from one table to the new table.

There is one other type of database security product. Rather than augmenting database security, however, a *security replacement* product supplants database security with an external package. This product's primary benefit is that it consolidates security administration for the data and the database. A second benefit is that the cascading REVOKE effect can be eliminated because most external data security packages do not cascade security revocations.

A security replacement product supplants database security with an external package.

Database security replacement tools have their weaknesses, too. These tools do not conform to the rigorous definition of the relational model, which states that the DBMS must control security. Some do not provide all types of database security. For example, system-level authorizations are frequently omitted. Another drawback is that if the external security package fails, the data is unprotected.

Data Warehousing, Analytics, and Business Intelligence

The data warehouse administrator requires a lot of help to automate the management of warehouse databases and to assist users in effectively querying the data warehouse.

ETL Tools

ETL stands for *extract, transform, and load*. ETL tools ease the burden of data movement and transformation because the tools understand the data format and environment of each DBMS they work with. A company's choice of ETL tools depends on the following factors:

- How many DBMS products need to be supported? If the project requires only a single DBMS (unlikely), the ETL features built into that DBMS might be sufficient for the project's needs.

- To what extent must data be replicated across multiple DBMS products?
- Does the data have to be synchronized across DBMS products?
- How much data manipulation and transformation is required as data moves from the OLTP source systems to the data warehouse?
- Is the data static or dynamic?
- If it is dynamic, is it updated online, in batch, or both?

The answers to these questions will help to determine the type of data warehousing tool necessary. Two basic types of data conversion tools are popular in the market today: replication and propagation.

Replication tools extract data from external application systems and heterogeneous databases for population into target database tables. Sometimes the target databases supported by this type of tool are narrow, but the source databases typically are broad. For example, a tool may be able to extract data from IMS, Oracle, SQL Server, non-database files, and other structures but will only support loading into a DB2 table. Before choosing a replication tool, be sure you know the full list of source and target data structures that must be supported, and choose your replication tool with care.

Choose replication tools with care.

Propagation tools capture data as it changes at the source (external applications and heterogeneous databases) and then insert it into target database tables. A propagation tool is similar in function to a replication tool, but propagation tools are active. They constantly capture updates made in the external system, either for immediate application to the target tables or for subsequent batch updating. This differs from the extract tool, which captures entire data structures, not data modifications.

In heterogeneous data-processing shops, the same data may need to be stored in multiple DBMS products. In a multiple-DBMS environment, the movement of data from DBMS to DBMS is a tedious task. ETL tools can be useful for converting data from DBMS to DBMS even outside of warehousing needs.

Query and Reporting Tools

Most DBMS products bundle a simple query tool with the product. Usually

this tool is very rudimentary and only permits SQL statements to be run with no formatting of the output. Most organizations will find these tools inadequate for developing professional, formatted reports or complete applications. These tools may also be inadequate for inexperienced users or those who want to develop or execute ad hoc queries.

Query and report-writing tools address these deficiencies. The capability to format reports without programming is probably their greatest asset. Another important feature is a GUI interface that allows you to develop data manipulation requests without writing complex SQL.

Query and report-writing tools address the deficiencies of bundled query tools.

Many vendors provide table query and reporting tools that can be used to enhance ad hoc query capabilities. These products typically

- Provide standard query formats and bundled reports
- Provide access to multiple file formats and heterogeneous databases
- Deliver extensive formatting options
- Offer more analytical functions than SQL
- Use static or dynamic SQL for stored queries
- Execute database commands from the query tool

Database query tools typically use point-and-click, drag-and-drop technology to greatly ease the report generation process. Fourth-generation languages (4GL) are also popular for accessing relational data. These products provide more functionality than a report-writing tool and use a GUI front end that makes them easier to use than 3GL programming languages such as Visual Basic, COBOL, or C. Generally, 4GL tools work in one of three ways:

- Queries are developed using 4GL syntax, which then is converted “behind the scenes” into SQL queries.
- SQL is embedded in the 4GL code and executed much like SQL embedded in 3GL code.

Fourth-generation languages are popular for accessing relational data.

- A hybrid methodology is used in which the executed SQL is either difficult or impossible to review.

Newer DBMS versions and releases now include more advanced support of queries, reports, and OLAP activities. For example, Microsoft SQL Server Analysis Services provides in-depth analytical query capabilities rivaling those of third-party OLAP tools. Of course, these services work for SQL Server only. Organizations that have to support analytical and query processing against heterogeneous databases should consider using a third-party tool that supports all of the DBMS products in the organization.

Advanced Analytics

Advanced analytics is a business-focused approach, comprising techniques that help build models and simulations to create scenarios, understand realities, and hypothesize future states. Advanced analytics uses data mining, predictive analytics, applied analytics, statistics, and other approaches in order to enable organizations to improve their business performance.

Traditional business intelligence enables us to understand the here and now, and even some of the why, of a given business situation. Advanced analytics goes deeper into the “why” of the situation and delivers likely outcomes. Although advanced analytics cannot infallibly predict the future, it can provide astute models for judging the likelihood of upcoming events. By allowing business managers to be aware of likely outcomes, advanced analytics can help to improve business decision making with an understanding of the effect those decisions may have in the next quarter or year.

Advanced analytical capabilities can be used to drive a wide range of applications, from tactical applications such as fraud detection to strategic analysis such as customer segmentation. Regardless of the applications, advanced analytics provides intelligence in the form of predictions, descriptions, scores, and profiles that help businesses better understand customer behavior and business trends.

Advanced analytics cannot infallibly predict the future, but it can provide astute models for judging the likelihood of upcoming events.

Most DBAs will not use advanced analytics tools but may be involved in

setting them up to work with the data warehouse and other databases.

Programming and Development Tools

Some tools straddle the line between DBA tool and programmer tool. For example, third-party vendors offer many tools that enhance the database application development environment. Programmers commonly use these tools, but DBAs quite frequently are involved in the review and analysis of such products. Some of the tasks that programming and development tools can perform include

- Testing SQL statements in a program editor while the programmer codes the SQL.
- Performing predictive performance estimation for programs and SQL statements.
- Explaining SQL statements in an editing session.
- Generating complete code from in-depth specifications. Some tools even generate SQL. When code generators are used, ensure that the generated code is efficient before promoting it to production status.
- Providing an enhanced environment for coding procedural SQL when creating triggers, stored procedures, and user-defined functions.
- Interfacing with 4GLs to create database access routines.

Because of the many varieties of database programming products, these tools should be evaluated on a case-by-case basis.

Checkpoint/Restart Tools

Some DBMS products provide no inherent capability for storing checkpoint information. When a program fails, a restart is difficult because there is no track of the last data that was committed to the database unless the programmer coded it into the program. Tools are available for some DBMS products to store checkpoint information that can be used by application programs during restarts. Automation of the restart process is another benefit provided by some checkpoint/restart solutions. Such products are useful for large batch database applications that issue many COMMITs. Many of these products require little, if any, modification of the program.

Testing Tools

Testing tools enable you to issue a battery of tests against a test base and

analyze the results. Testing tools typically are used for all types of applications, but some have been specifically extended to support testing against database tables.

Debugging Tools

SQL debugging tools are useful for identifying and fixing syntax and logic errors in difficult SQL statements. These tools are most useful for procedural SQL dialects such as PL/SQL and Transact-SQL.

Miscellaneous Tools

In addition to the tools we have reviewed already, there are all kinds of niche tools available. Some of them are non-database-specific tools, while others are simply useful single-purpose programs that help when managing database systems. Let's examine some of these miscellaneous tool categories.

All kinds of niche tools are available.

Space Management Tools

Most DBMSs provide basic statistics for space utilization, but the in-depth statistics required for both space management and performance tuning are usually inadequate for heavy-duty administration. For example, most DBMSs lack the ability to monitor the requirements of the underlying files used by the DBMS. When these files go into extents or become defragmented, performance can suffer. Without a space management tool, the only way to monitor this information is with arcane and difficult-to-use operating system commands. This can be a tedious exercise.

Additionally, each DBMS allocates space differently. The manner in which the DBMS allocates space can result in inefficient disk usage. Sometimes space is allocated, but the database will not use it. A space management tool is the only answer for ferreting out the amount of used space versus the amount of allocated space.

Each DBMS allocates space differently.

Space management tools often interface with other database and systems management tools such as operating system space management tools, database analysis tools, system catalog query and management tools, and database utility generators.

Compression Tools

A standard tool for reducing storage costs is the compression utility. This type of tool operates by applying an algorithm to the data in a table such that the data is encoded in a more compact area. When the amount of area needed to store data is reduced, overall storage costs are decreased. Compression tools must compress the data when it is added to the table and subsequently modified, then expand the data when it is later retrieved. Additionally, some tools are available that compress database logs, enabling more log information to be retained on disk before it is off-loaded to another medium.

Third-party compression routines are usually specified by using APIs at the database or operating system level. Sometimes tables must be dropped and recreated to apply the compression routine.

In general, a compression algorithm increases CPU costs while it decreases storage and I/O costs. This trade-off is not beneficial for all tables. For example, if a compression routine saves 30 percent on storage costs but increases CPU consumption without decreasing I/O, the trade-off is probably not beneficial.

A compression tool can decrease storage by reducing the size of the rows to be stored. CPU consumption usually increases because additional processing is required to compress and expand the row. I/O costs, however, could decrease—perhaps dramatically so.

A compression tool can reduce the size of the rows to be stored.

Some DBMSs support compression without requiring a third-party product. The major advantage of third-party compression tools is that most of the vendors provide additional flexibility and multiple compression algorithms for different types of data.

Online Standards Manuals

Products are available that provide “canned” standards for implementing, accessing, and administering databases. These tools are particularly useful for new database shops. By purchasing an online standards manual, new—and even more experienced—shops can quickly come up to speed with each new DBMS.

Everyone benefits from these products if the third-party vendor

automatically ships updates whenever there is a new version or release of the DBMS product. Online manuals can function as cheap training in the new DBMS release (e.g., migrating from Oracle8i to Oracle9i).

Online manuals can function as cheap training.

Any product that delivers database standards should

- Provide online networked access so that all developers and DBAs can access the manual
- Be extensible, so additional standards can be added
- Be modifiable, so the provided standards can be altered to suit prior shop standards (e.g., naming conventions, programming standards)

Other Tools

Many types of database tools are available. The categories in this chapter cover the major types of tools, but not all tools can be easily pigeonholed. For example, consider a database space calculator. It reads table DDL and information on the number of rows in the table to estimate space requirements. A space calculator is often provided with another tool such as a space management tool or a database design and modeling tool.

Additionally, some standard system management and programming tools can be used to optimize the database environment. For example, a third-party sorting tool can be used to presort data before loading or to sort output results for comparison during database program testing.

Furthermore, new types of products are announced quite regularly. For example, one vendor recently released a product that analyzes the SQL that accesses a particular table and uses a set of expert rules to recommend the most efficient indexes. Be sure to keep your eyes open for new and useful types of tools not mentioned in this chapter.

New types of products are announced quite regularly.

Examine Native DBA Tools

Before moving forward to purchase any third-party DBA tools, be sure to thoroughly examine the functionality provided by the DBMS vendor. DBMS vendors have greatly improved the features offered in their native tools, and

if you have not investigated them recently, you should take another look.

For example, Oracle Enterprise Manager and Microsoft SQL Server SSMS have become very feature-rich products. Be careful, though, as sometimes only a base set of functionality is provided out of the box with additional features requiring additional licensing and, importantly, additional cost.

Homogeneous shops, with only a single DBMS to support, frequently can do well using the DBMS vendor's DBA tools. Heterogeneous shops, with multiple DBMS products, may do better by purchasing a tool from a third party that manages multiple database systems.¹

Evaluating DBA Tool Vendors

Advanced organizations augment their DBA staff with advanced software tools that help to automate the task of database administration. Multiple categories of database tools are available to organizations to improve the efficiency of database administration. Sometimes the DBA will have to choose from a list of similar tools provided by multiple vendors or investigate a new vendor that offers a tool not available from the primary vendor.

Although the most important aspect of database tool selection is the functionality of the tool and the way it satisfies the needs of your organization, the nature and stability of the product vendor are important also. This section provides suggested questions to ask when you are selecting a database tool vendor.

The nature and stability of the product vendor are important.

- How long has the vendor been in business? How long has the vendor been supplying tools for the particular DBMS products you manage? In general, the longer the vendor has been in business, the better; and the longer the vendor has worked with the DBMS products in question, the better.
- Does your company have other DBA or system management tools from this vendor? How satisfied are the users of those tools?
- Are other organizations satisfied with the tool you are selecting? Obtain a list of other companies that use the same tool, and contact several of them.
- Does the vendor provide a 24-hour support number? If not, what are the

vendor's hours of operation? Does the vendor have a toll-free number? If not, be sure to determine the cost of long-distance charges when you are requesting customer support. You want to avoid accumulating significant long-distance charges when you are requesting customer support from a vendor. (You may be able to use Skype or another VOIP technology to minimize these costs.)

- Does the vendor provide online technical support? Can you access it to evaluate its usefulness before establishing a relationship with the vendor? If so, scan some of the FAQs and reported problems for the tools before committing to the vendor's product. Keep in mind that the Web support should be in addition to phone support, not a replacement.
- Evaluate the response of the technical support number. Call the number with technical questions at least five times throughout the day: before 8:00 A.M., around noon, just before 5:00 P.M., after 9:00 P.M., and just after midnight. These are the times when you could find problems with the level of support provided by the vendor. Was the phone busy? Were you put on hold? For how long? When you got a response, was it accurate and friendly? Did the person who answered the phone have to find someone with more technical knowledge? (This can indicate potential problems.) If you really want to test the vendor, call after midnight to gauge the response to an after-hours support request.
- Does this vendor supply other tools that your organization might need later? If so, are they functionally integrated with this one? Does the vendor supply a full suite of products or just a few? Does the vendor provide solutions in each of the categories discussed above, or in just a few? It is wise to work with a vendor that can satisfy all (or most) of your database administration needs, not just a few of them. However, a multivendor strategy can be helpful to minimize cost (because you can pit one vendor against another to negotiate discounts).
- Are the vendor's tools integrated with other tools? For example, a product that analyzes databases to determine whether a reorg is required should integrate the reorg with your shop's job scheduler.
- Does the vendor provide training? Is it on-site training? Does the vendor supply DBMS training as well as training for its tools? Is computer-based training available for the vendor's tools?
- Are installation, technical, and user manuals available in both hard and

soft copy? Are hard copies provided free of charge? If so, how many copies? If not, what is the charge?

- Will the vendor answer database questions free of charge in addition to questions about its product? For example, if the vendor provides Oracle tools, are the support personnel skilled enough to answer a question about Oracle Database administration? Sometimes vendors will do this, but they do not typically advertise the fact. Try it out by calling the technical support number.
- Does the vendor have a local office? If not, are technicians readily available for on-site error resolution if needed? If so, at what price?
- Will the vendor deliver additional documentation or error resolution information by overnight mail? Does the vendor publish a fax number? An e-mail address?
- How are software fixes provided? Electronically? By tape? Over the Web? Via FTP? Is a complete reinstallation required? Are fixes typically accomplished using zaps? For mainframe folks, does the vendor support SMP/E?
- How many man-hours, on short notice, is the vendor willing to spend to solve problems? Is there a guaranteed time limit to resolve severe product problems?
- Is the vendor willing to send a sales representative to your site to do a presentation of the product tailored to your needs? Or is it a “phone sales only” organization? If the vendor will not send someone on-site, it can indicate that you are dealing with a small vendor or a vendor that could be difficult to deal with later.
- Is the vendor a business partner of DBMS vendors such as Microsoft, IBM, and Oracle? How soon will the vendor’s tools be modified to support new releases and versions of the DBMS?
- Have the vendor’s tools been reviewed or highlighted in any industry publications recently? If so, obtain the publications and read the articles. Sometimes the articles will be available on the Web. Make sure that they were not written by the vendor, but by a neutral third party—preferably a customer, consultant, or industry analyst.
- Will the vendor supply a list of consultants who have used the product and can vouch for its functionality?

- Will the vendor assist in developing a cost justification? Most tool vendors are eager for your business and will be more than willing to provide cost justifications to help you sell upper management on the need for DBA tools.
- Does the vendor provide sample scripts (or JCL for the mainframe) to run its product? Is an easy-to-use “Getting Started” manual available?
- Does the vendor charge an upgrade fee when you upgrade your hardware or DBMS? How flexible are the terms and conditions of the contract?
- Is the vendor willing to set a ceiling for increases in the annual maintenance charge?
- If the vendor is sold or goes out of business, will the vendor supply the source code of the tool? Is the source code escrowed? If not, are the terms and conditions of the contract flexible in the event the vendor is acquired? Given the state of the industry today, with mass vendor consolidation, this is an important item to consider or you may be stuck with unsupported products or a “difficult” vendor post acquisition.
- What is the current status of the vendor? Have any recent events or business downturns resulted in lower market share? Did the company recently undergo a layoff? Are there any outstanding lawsuits against the company? Answers to these questions can help to identify the stability of the vendor.
- Does the vendor supply database administration tools for all of the DBMS products used at your shop? Can the same tool, using the same interface, be used to manage multiple databases across multiple operating systems?
- How does the vendor rank enhancement requests? Is every request ranked equally or are certain customers favored over others? If so, how can you be assured that you will be one of the vendor’s “favored” customers?

These questions can provide the basis for your organization’s evaluation of DBA tool vendors. Judge for yourself which criteria are most important to your organization. You might consider applying a weighting factor to each question to rate the responses. If you do so, be sure not to inform the vendor of this weighting to ensure that any help the vendor may be supplying you is

“honest” help.

Homegrown DBA Tools

Some organizations build homegrown DBA tools to help support their DBMS environments. Such tools are usually developed and maintained by the DBAs or systems programmers. I have seen homegrown tools that address just about every aspect of database administration that vendor-marketed tools address.

There are pros and cons to pursuing a homegrown DBA tool strategy. Most organizations view the biggest benefit as decreased cost. This may be true in some cases, but it often is not. Although homegrown tools will not incur the initial expense of a third-party tool, the ongoing maintenance expense can be prohibitive. Every time a new version or release of the DBMS is installed, a homegrown tool will need to be at least retested and usually recoded. Not only does this incur a cost, but it takes valuable management time away from the DBA. Instead of managing the organization’s databases, the DBA must maintain the homegrown tools.

Homegrown tools have pros and cons.

New DBAs need to understand the DBA team’s tool philosophy when coming on board. Does the company build its own tools or purchase third-party tools? Sometimes the answer is a mixture of the two approaches.

Summary

Third-party DBA tools can significantly improve the efficiency of applications that access relational data. When evaluating products, look for features that are important to your organization. Consider adopting checklists for product comparisons based on the features discussed in this chapter. Remember, most DBMSs leave quite a bit to be desired in the administration, data access, performance monitoring, and application development areas. Third-party DBA tools can help to minimize these weaknesses and alleviate your database administration pains.

Adopt checklists for product comparisons.

Review

1. Name the three categories of performance tools and briefly summarize their typical capabilities.
2. If you are asked by management to develop a cost justification for DBA tools, how can you minimize the amount of work it takes to produce such a justification?
3. What type(s) of tool(s) would be most helpful to an organization that is brand-new to database management, and why?
4. What is the purpose of a catalog visibility tool?
5. What is the benefit of purchasing database utilities from a third-party vendor when most DBMS products already ship with a set of utilities “in the box”?
6. What type of problem can be caused when using a table editor to modify data in a production database table?
7. What does the acronym [*ETL*](#) mean?
8. Why would it be important to contractually guarantee that the vendor escrow the source code for any tools your organization purchases?
9. What is the difference between a propagation tool and a replication tool?
10. Name the seven questions that a database auditing tool should be able to answer.

24. DBA Rules of Thumb

The first 23 chapters of this book cover the practices, policies, and procedures associated with the management discipline of database administration. This chapter provides guidance for DBAs as they embark on the job of implementing those practices, policies, and procedures. Database administration is a very technical discipline; it is also a very political discipline—the DBA is highly visible within the organization. DBAs should be armed with a proper attitude and sufficient technical knowledge before attempting to practice database administration.

DBAs should be armed with a proper attitude and sufficient technical knowledge.

The following sections outline some basic guidelines to help you become a well-rounded, respected, and professional DBA.

Write Down Everything

During the course of performing your job as a DBA, you are likely to encounter many challenging tasks and time-consuming problems. Be sure to document the processes you use to resolve problems and overcome challenges. Such documentation can be very valuable should you encounter the same, or a similar, problem in the future. It is better to read your notes than to try to recreate a scenario from memory.

Written notes and details also can be shared with your peers as a learning aid. Notes that document the steps taken to overcome problems can be used to improve subsequent DBA efforts and may prevent problems in the future.

Keep Everything

DBA is just the job for you if you are a pack rat. It is a good practice to keep everything you come across during the course of your job. This includes, but is not limited to,

Be a pack rat.

- Manuals—old and new, online and printed

- Scripts—to make changes, run reorgs, whatever
- Programs—quick-fix programs, sample programs, and so on
- Proceedings and presentations—from webinars, user groups, and conferences
- Notes—from design reviews, specification meetings, team meetings, user groups and conferences, and so on
- Textbooks—from any classes you have attended (as well as additional technical books such as this one)
- White papers
- Reports—from performance monitors, queries, applications, and so on
- Articles—from magazines or links to Web sites
- Reference manuals—for all of your DBMS products and tools
- CDs—containing product demos, products, white papers, presentations, and so on

Of course, keeping all of that information organized and accessible can be a difficult challenge, although keeping everything from processes, to links, to scripts, and so on can and will help DBAs to determine a proper course of action when they find themselves in a jam. But organization, foresight, and preparation are also required.

Being prepared for “whatever may happen” is not about having an advanced set of DBA tools or a bunch of archived information, but having a sound understanding of your DBMS, databases, supporting infrastructure, and the interrelationship of the business with your data, and working from a plan. Waiting for a failure to occur, and only then scrambling to find that “perfect” script (that you may never have even tested), is a recipe for failure.

Automate!

Why should you do it by hand if you can automate? Anything you do can probably be done better by a computer—if it is programmed properly. Once a task is automated, you save yourself valuable time that probably will be spent tackling other problems.

Furthermore, don’t reinvent the wheel. Someone, somewhere, at some time may have already solved the problem you are attempting to solve. Look for Web sites that allow you to download and share scripts. If you have the

budget, purchase DBA tools as discussed in [Chapter 23](#).

Don't reinvent the wheel.

Of course, you can take the automation idea too far. There has been a lot of talk and vendor hype about self-managing database systems. And DBMS vendors have made great strides in terms of automated administration features. But even so, there are many reasons why DBAs are not on the fast path to extinction. Self-managing databases systems, although a laudable goal, are not close to delivering a “lights-out” DBMS environment. Little by little and step by step, database maintenance and performance management are being improved, simplified, and automated. But you know what? DBAs will not be automated out of existence in my lifetime—and probably not in your children's lifetime either.

Many of the self-managing features require using the built-in (or freely available) tools from the DBMS vendor, such as Oracle Enterprise Manager or IBM Data Studio. But many organizations prefer to use heterogeneous solutions that can administer databases from multiple vendors (Oracle, IBM, Microsoft, etc.), all from a single console. Most of these tools have had self-managing features for years and yet they have not made the DBA obsolete.

And let's face it, a lot of the DBMS vendors' claims are more hyperbole than fact. Some self-managing features are announced years before they will become generally available in the DBMS. All vendors' claims to the contrary, no database today is truly 100 percent self-contained. Every database needs some amount of DBA management—even when today's best self-management features are being used.

What about the future? Well, things will get better—and probably more costly. You don't think the DBMS vendors are building this self-management technology for free, do you? But let's remove cost from the equation for a moment. What can a self-managing database actually manage?

Most performance management solutions allow you to set performance thresholds. A threshold allows you to set up a trigger that says something like “When x percent of a table's pages contain chained rows or fragmentation, schedule a reorganization.” But these thresholds are only as good as the variables you set and the actions you define to be taken upon tripping the threshold. Some software is intelligent; that is, it “knows” what to do and when to do it. Furthermore, it may be able to learn from past actions and

results. The more intelligence that can be built into a self-managing system, the better the results typically will be. But who among us currently trusts software to work like a grizzled veteran DBA? The management software should be configurable such that it alerts the DBA as to what action it wants to take. The DBA can review the action and give a “thumbs up” or “thumbs down” before the corrective measure is applied. In this way, the software can earn the DBA’s respect and trust. When the DBA trusts the software, he can turn it on so that it self-manages “on the fly” without DBA intervention. But today, in most cases, a DBA is required to set up the thresholds, as well as to ensure their ongoing viability.

Of course, not all DBA duties can be self-managed by software. Most self-management claims are made for performance management, but what about change management? The DBMS cannot somehow read the mind of its user and add a new column or index, or change a data type or length. This nontrivial activity requires a skilled DBA to analyze the database structures, develop the modifications, and deploy the proper scripts or tools to implement the change. Of course, software can help simplify the process, but software cannot replace the DBA.

Not all DBA duties can be self-managed by software.

Furthermore, database backup and recovery will need to be guided by the trained eye of a DBA. Perhaps the DBMS can become savvy enough to schedule a backup when a system process occurs that requires it. Maybe the DBMS of the future will automatically schedule a backup when enough data changes. But sometimes backups are made for other reasons: to propagate changes from one system to another, to build test beds as part of program testing, and so on. A skilled professional is needed to build the proper backup scripts, run them appropriately, and test the backup files for accuracy. And what about recovery? How can a damaged database know it needs to be recovered? Because the database is damaged, any self-managed recovery it might attempt is automatically rendered suspect. Here again, we need the wisdom and knowledge of the DBA. The DBA may rely upon an automation tool to simplify the recovery process, but the tool itself cannot perform the task without the guidance of an experienced DBA.

And there are many other DBA duties that cannot be completely automated. Because each company is different, the DBMS must be

customized using configuration parameters. Of course, you can opt to use the DBMS “as is,” right out of the box. But knowledgeable DBAs can configure the DBMS so that it runs appropriately for their organization. Problem diagnosis is another tricky subject. Not every problem is readily solvable by developers using just the “Messages and Codes” manual and a help desk. What happens with particularly thorny problems if the DBA is not around to help?

All of these caveats aside, automation should be a desired component of the database administration equation. Automate the routine tasks to free up more time to spend on the more difficult tasks and those that cannot be completely automated.

Share Your Knowledge

Knowledge transfer is an important part of being a good DBA—both transferring your knowledge to others and participating in having others’ knowledge transferred to you.

The more you learn, the more you should share that knowledge with others. Many local database user groups meet quarterly or monthly to discuss aspects of database management systems. Be sure to attend these sessions to learn what your peers are doing. When you have some good experiences to share, put together a presentation yourself. Sometimes you can learn far more by presenting at these events than by simply attending because the attendees will seek you out to discuss your approach. You might also use local user groups as a springboard for presenting at international database user group conferences hosted by the major user group associations (see the sidebar “[Major Database User Group Associations](#)”).

Major Database User Group Associations

User group associations exist for each of the major DBMSs. These associations host regional meetings and international conferences and provide resources for database users and DBAs.

- IDUG: International DB2 User Group www.idug.org
- IIUG: International Informix Users Group www.iiug.org
- IOUG: Independent Oracle Users Group www.ioug.org
- ISUG: International Sybase User Group www.isug.com

- Professional Association for SQL Server www.sqlpass.org

Be sure to join the association for the DBMSs that you use and participate in the opportunities provided for networking and career advancement.

Another avenue for sharing your knowledge is using one of the many online forums. Web portals and Web-based publications are constantly seeking out content for their Web sites. Working to put together a tip or article for these sites helps you to organize your thoughts and to document your experiences. Furthermore, as the author of an article or column on the Web, you will gain exposure to your peers. Sometimes this type of exposure can help you to land that coveted job.

Use online forums.

You also should consider participating in social networking sites such as Facebook, LinkedIn, and Twitter. Each of these social network communities can be used to further your career and add to your professional network. Facebook is useful for more casual interaction, but many database professional participate, and the community developed through Facebook connections can be beneficial for broad interaction. LinkedIn is a more professional network and is useful for keeping in touch with coworkers and contacts as they move from job to job. Twitter is probably the best method for interacting with other DBAs and sharing stories, advice, and more. For details refer to the sidebar “[Twitter](#).”

Twitter

The basic idea of Twitter (www.twitter.com) is simple. It offers a platform for users to publish messages of no more than 140 characters at a time. And that can seem limiting . . . until you’ve used Twitter for a while. When you tweet—that is, send a Twitter message—it gets broadcast to all of your followers. Likewise, when anyone you follow sends a tweet, it gets broadcast to you.

For example, you can follow me (the author) on Twitter at [@craigmullins](#). I share information on my blogs, writings, speaking engagements, and so on. A well-developed feed of folks you follow on Twitter can be a very useful conduit for news, trends, issues, and ideas.

Many professional DBAs participate on Twitter, as well as IT luminaries such as Bill Gates (@BillGates), Robert Scoble (@scobleizer), and Don Tapscott (@dtapscott). You can also follow industry analysts using Twitter, including analysts from Red Monk, Gartner, IDC, and Forrester.

You can also send direct messages to other users who follow you. This can reduce the use of e-mail for short communiqués that do not require in-depth follow-up or attachments. Some forward-thinking companies have Twitter customer advocates who keep a watchful eye on the tweetstream and respond to complaints, comments, and suggestions.

When participating in Twitter, be sure to use hashtags. A hashtag is a term preceded by the # character, such as #DBA. The hashtag acts as metadata to inform other Twitter users about the content of your message.

Using Twitter, you can expand your network and interact with DBAs all over the world.

Finally, if you have the time, consider writing for one of the database-related magazines. Writing an article for publication will take the most amount of time, but it can bring the most exposure. Furthermore, some journals will pay you for your material.

However, the best reason to share your knowledge with others is because you want them to share their knowledge and experiences with you. A spirit of cooperation maintains the community of DBAs who are willing and eager to provide assistance.

Analyze, Simplify, and Focus

As we learned in the previous chapters, the job of a DBA is complex and spans many diverse technological and functional areas. It is easy for a DBA to get overwhelmed with certain tasks—especially those that are not performed regularly. In a complex, heterogeneous, distributed world it can be hard to keep your eye on the right ball, at the right time. The best advice I can give you is to remain focused and keep a clear head.

Understand the purpose for each task and focus on performing the steps that will help you to achieve that end. Do not be persuaded to broaden the

scope of work for individual tasks unless it cannot be avoided. In other words, don't try to boil the ocean. If nonrelated goals get grouped together into a task, it can become easy to work long hours with no clear end in sight.

I am not saying that a DBA should (necessarily) specialize in one particular area (e.g., performance). What I am suggesting is that each task should be given the appropriate level of focus and attention to details. Of course, I am not suggesting that you should not multitask either. The successful DBA will be able to multitask while giving full attention to each task as it is being worked on.

What is the enemy of focus? There are many: distraction, lack of knowledge, "management," and always worrying about the next thing to try or do. Such distractions can wreak havoc on tasks that require forethought and attention to detail.

Analyze, simplify, and focus. Only then will tasks become measurable and easier to achieve.

Don't Panic!

A calm disposition and the ability to remain cool under strenuous conditions are essential to the makeup of a good DBA. Problems will occur—nothing you can do can eliminate every possible problem or error. Part of your job as a DBA is to be able to react to problems with a calm demeanor and analytical disposition.

When a database is down and applications are unavailable, your environment will become hectic and frazzled. The best things you can do when problems occur are to remain calm and draw on your extensive knowledge and training. As the DBA, you will be the focus of the company (or at least the business units affected) until the database and applications are brought back online. It can be a harrowing experience to recover a database with your boss and your users hovering behind your computer terminal and looking over your shoulder. Be prepared for such events, because eventually they will happen. Panic can cause manual errors—the last thing you want to happen when you are trying to recover from an error.

Panic can cause manual errors.

The more comprehensive your planning and the better your procedures, the

faster you will be able to resolve problems. Furthermore, if you are sure of your procedures, you will remain much calmer.

Measure Twice, Cut Once

Being prepared means analyzing, documenting, and testing your DBA policies and procedures. Creating procedures in a vacuum without testing will do little to help you run an efficient database environment. Moreover, it will not prepare you to react rapidly and effectively to problem situations.

The old maxim applies: Measure twice, cut once. In the case of DBA procedures, this means analyze, test, and then apply. Analyze your environment and the business needs of the databases to create procedures and policies that match those needs. Test those procedures. Finally, apply them to the production databases.

Understand the Business, Not Just the Technology

Remember that being technologically adept is just a part of being a good DBA. Although technology is important, understanding your business needs is more important. If you do not understand the impact on the business of the databases you manage, you will simply be throwing technology around with no clear purpose.

Business needs must dictate what technology is applied to what database—and to which applications. Using the latest and greatest (and most expensive) technology and software might be fun and technologically challenging, but it most likely will not be required for every database you implement. The DBA's tools and utilities need to be tied to business strategies and initiatives. In this way, the DBA's work becomes integrated with the goals and operations of the organization.

The first step in achieving this needed synergy is the integration of DBA services with the other core components of the IT infrastructure. Of course, DBAs should be able to monitor and control the databases under their purview, but they should also be able to monitor them within the context of the broader spectrum of the IT infrastructure—including systems, applications, storage, and networks. Only then can companies begin to tie service-level agreements to business needs, rather than technology metrics.

DBAs should be able to monitor and control the databases

under their purview.

DBAs should be able to gain insight into the natural cycles of the business just by performing their job. Developers and administrators of other parts of the IT infrastructure will not have the vision into the busiest times of the day, week, quarter, or year because they are not privy to the actual flow of data from corporate business functions. But the DBA has access to that information as a component of performing the job. It is empowering to be able to understand business cycle information and apply it on the job.

DBAs need to expand further to take advantage of their special position in the infrastructure. Talk to the end users—not just the application developers. Get a sound understanding of how the databases will be used before implementing any database design. Gain an understanding of the database's impact on the company's bottom line, so that when the inevitable problems occur in production you will remember the actual business impact of not having that data available. This also allows you to create procedures that minimize the potential for such problems.

To fulfill the promise of business/IT integration, it will be necessary to link business services to the underlying technology. For example, a technician should be able to immediately comprehend that a service outage to transaction X7R2 in the PRD2 environment means that regional demand deposit customers cannot access their accounts. See the difference?

Focusing on transactions, TP monitors, and databases is the core of the DBA's job. But servicing customers is the reason the DBA builds those databases and manages those transactions. Technicians with an understanding of the business impact of technology decisions will do a better job of servicing the business strategy. This is doubly true for the DBA's manager. Technology managers who speak in business terms are more valuable to their company.

Of course, the devil is in the details. A key component of realizing effective business/IT integration for DBAs is the ability to link specific pieces of technology to specific business services. This requires a service impact management capability—that is, analyzing the technology required to power each critical business service and documenting the link. Technologies exist to automate some of this through event automation and service modeling. Such capabilities help to transform availability and performance

data into detailed knowledge about the status of business services and service-level agreements.

Today's modern corporation needs technicians who are cognizant of the business impact of their management decisions. As such, DBAs need to get busy transforming themselves to become more business savvy—that is, to keep an eye on the business impact of the technology under their span of control.

Corporations need technicians who are cognizant of the business impact of their management decisions.

Don't Become a Hermit

Work with the application developers. Don't isolate yourself in your own little DBA corner of the world. The more you learn about what the applications do and the application requirements, the better you can adjust and tune the databases to support those applications.

A DBA should be accessible. Don't be one of those DBAs whom everyone is afraid to approach. The more you are valued for your expertise and availability, the more valuable you are to your company.

Be accessible.

Use All of the Resources at Your Disposal

Remember that you do not have to do everything yourself. Use the resources at your disposal. We have talked about some of those resources, such as articles and books, Web sites and scripts, user groups and conferences. But there are others.

Do not continue to struggle with problems when you are completely stumped. Some DBAs harbor the notion that they have to resolve every issue themselves in order to be successful. Sometimes you just need to know where to go to get help to solve the problem. Use the DBMS vendor's technical support, as well as the technical support line of your DBA tool vendors. Consult internal resources for areas where you have limited experience, such as network specialists for network and connectivity problems, system administrators for operating system and system software problems, and security administrators for authorization and protection issues.

Know where to go to get help to solve the problem.

When you go to user groups, build a network of DBA colleagues whom you can contact for assistance. Many times others have already encountered and solved the problem that vexes you. A network of DBAs to call on can be an invaluable resource (and no one at your company even needs to know that you didn't solve the problem yourself).

Finally, be sure to understand the resources available from your DBMS vendors. DBMS vendors offer their customers access to a tremendous amount of useful information. All of the DBMS vendors offer software support on their Web sites. Many of them provide a database that users can search to find answers to database problems. IBM customers can use IBMLink,¹ and both Oracle and Microsoft offer a searchable database in the support section of their Web sites. Some DBAs claim to be able to solve 95 percent or more of their problems by researching online databases. These resources can shrink the amount of time required to fix problems, especially if your DBMS vendor has a reputation of “taking forever” to respond to issues.

Of course, every DBA should also be equipped with the DBMS vendor's technical support phone number for those tough-to-solve problems. Some support is offered on a pay-per-call basis, whereas other times there is a prepaid support contract. Be sure you know how your company pays for support before calling the DBMS vendor. Failure to know this can result in your incurring significant support charges.

Keep Up-to-Date

Keep up-to-date on technology. Be sure that you are aware of all of the features and functions available in the DBMSs in use at your site—at least at a high level, but preferably in depth. Read the vendor literature on future releases as it becomes available to prepare for new functionality before you install and migrate to new DBMS releases. The sooner you know about new bells and whistles, the better equipped you will be to prepare new procedures and adopt new policies to support the new features.

Keep up-to-date on technology in general, too. For example, DBAs should understand new data-related technologies such as NoSQL, but also other newer technologies that interact with database systems. Don't ignore industry and technology trends simply because you cannot immediately think of a

database-related impact. Many non-database-related “things” (for example, XML) eventually find their way into DBMS software and database applications.

Keep up-to-date on industry standards—particularly those that impact database technology such as the SQL standard. Understanding these standards before the new features they engender have been incorporated into your DBMS will give you an edge in their management. DBMS vendors try to support industry standards, and many features find their way into the DBMS because of their adoption of an industry standard.

Keep up-to-date on technology in general and on industry standards.

As already discussed, one way of keeping up-to-date is by attending local and national user groups. The presentations delivered at these forums provide useful education. Even more important, though, is the chance to network with other DBAs to share experiences and learn from each other’s projects.

Attend local and national database user groups.

Through judicious use of the Internet and the Web, it is easier than ever before for DBAs to keep up-to-date. Dozens of useful and informative Web sites provide discussion forums, script libraries, articles, manuals, and how-to documents. Consult [Appendix D](#) for a list of DBA-focused Web resources. Remember, though, this is just a starting point.

Invest in Yourself

Most IT professionals continually look for their company to invest money in their ongoing education. Who among us does not want to learn something new—on company time and with the company’s money? Unless you are self-employed, that is!

Yes, your company should invest some funds to train you on new technology and new capabilities, especially if it is asking you to do new things. And since technology changes so fast, most everyone has to learn something new at some point every year. But the entire burden of learning should not be placed on your company.

Budget some of your own money to invest in your career. After all, you

probably won't be working for the same company your entire career. Why should your company be forced to bankroll your entire ongoing education? Now, I know, a lot depends on your particular circumstances. Sometimes we accept a lower salary than we think we are worth because of the "perks" that are offered. And one of those perks can be training.

Budget some of your own money to invest in your career.

But some folks simply abhor spending *any* of their hard-earned dollars to help advance their careers. Shelling out a couple of bucks to buy some new books, subscribe to a publication, or join a professional organization shouldn't be out of the reach of most DBAs, though.

A willingness to spend some money to stay abreast of technology is a trait that should apply to DBAs. Most DBAs are insatiably curious, and many are willing to invest some of their money to learn something new. Maybe they bought that book on XML before anyone at their company started using it. Perhaps it is just that enviable bookshelf full of useful database books in their cubicle. Or maybe they paid that nominal fee to subscribe to the members-only content of that SQL Server portal. They could even have forked over the \$10 fee to attend the local user group.

Don't get me wrong. I'm not saying that companies should not reimburse for such expenses. They should, because it provides for better-rounded, more educated, and more useful employees. But if your employer won't pay for something that you think will help your career, why not just buy it yourself?

And be sure to keep a record of such purchases because unreimbursed business expenses can be tax deductible.

Summary

The job of the DBA is a challenging one—from both a technological and a political perspective. Follow the rules of thumb presented in this chapter to increase your likelihood of achieving success as a DBA.

Final Exam

1. At a high level, discuss the DBA's primary job responsibilities.
2. Is a certified DBA necessarily a qualified DBA? Why or why not?
3. Why must the DBA be prepared to function as a jack-of-all-trades?

4. Why should database standards be implemented, and what are the risks associated with a lack of DBA standards and procedures?
5. Name the TPC benchmarks and describe how they differ from one another.
6. Name five requirements that must be planned for when installing a new DBMS.
7. What is the difference between a conceptual data model and a logical data model?
8. Why is data modeling important for database development?
9. Describe in broad terms the goals of normalization.
10. What is the difference between an entity and an entity occurrence?
11. Every attribute will do one of three things. Name those things and describe them.
12. What is the only reason to denormalize a physical data model?
13. Under what circumstances should you consider a bitmap index rather than a b-tree index?
14. Describe how a relational database uses indexes.
15. Why might the order in which columns are created in a table be important for physical database design?
16. If indexes are beneficial to performance, why not create every conceivable index just to be on the safe side?
17. Describe what the acronym *ACID* means, and define each component.
18. Why does locking assure data integrity?
19. What does *relational closure* mean and what is its significance in application design?
20. What is the difference between a lock time-out and a deadlock?
21. Why is it important to review application code in addition to reviewing SQL?
22. Cite several reasons for including representatives from application development management in design reviews.
23. Why should the DBA lead most of the design reviews?
24. Why is it difficult to add a column between two existing columns in a

table?

25. Explain what is meant by *cascading drop*.
26. If you have to drop an entire database to effect a database change, what other database structures will also be dropped automatically by the DBMS?
27. Define what is meant by *availability*.
28. What is the typical cause of an unplanned outage?
29. What are nondisruptive database utilities, and why are they important for maintaining database availability?
30. Define *database performance*.
31. What are the three “things” that need to be addressed in database application tuning?
32. What does the 80/20 rule mean, and how should it be applied to database application performance tuning?
33. According to most experts, what is the primary cause of performance problems in database applications?
34. How does the data cache (or buffer pool) improve the performance of database processing?
35. Why are certain database operations not logged?
36. What benefits can accrue by caching optimized SQL in memory?
37. What type of information is recorded on the database transaction log?
38. What is the greatest performance tuning technique a DBA can use to improve database performance?
39. What is the benefit of clustering data?
40. What are the causes of database and index disorganization?
41. What is the benefit of allocating tablespaces and indexes to separate disk devices?
42. Define what is meant by *physical data independence*.
43. What factors influence the optimizer during SQL access path selection?
44. Name and describe the two predominant relational join methods.
45. Under what circumstances will a table scan outperform indexed

access?

46. Define what is meant by *entity integrity*.
47. Describe the difference between database structural integrity and semantic data integrity.
48. What is the difference between a primary key constraint and a unique constraint?
49. Name and describe four types of database structural integrity problems that may be encountered by a DBA.
50. Explain two ways that views can be used to implement data security.
51. What two SQL DCL statements are used to establish and remove database privileges?
52. What is a cascading REVOKE and what effect can it have on database security?
53. What is the effect of granting a privilege to PUBLIC?
54. Why is it important to protect some database resources with security mechanisms external to the DBMS?
55. Describe how SQL injection attacks work.
56. Explain the purpose for creating an audit trail of database operations.
57. Describe data masking and explain how it benefits regulatory compliance efforts.
58. Why is a quiesce point important for a point-in-time recovery?
59. What factors determine whether a full or incremental image copy backup should be taken for a database object?
60. Name and describe the three types of database failures that may require recovery.
61. Name the different types of recovery, and discuss the factors that influence when each type of recovery should be performed.
62. Describe what is meant by *disaster* in terms of database administration.
63. What factors should be considered when determining the criticality of a database object for disaster recovery planning?
64. Once the disaster recovery plan is written, it should never be changed: true or false?

65. Why are most database files stored on disk devices instead of other storage media?
66. Your organization has decided to purchase RAID arrays to support your DBMS. Which levels of RAID would you advise them to use and why?
67. What advantages does a SAN have over SCSI devices?
68. Calculating the amount of storage required for a table is as simple as adding up the number of bytes for the columns in a row and multiplying the sum by the total number of rows: true or false?
69. Why is a two-phase COMMIT necessary when data is modified at two locations within a single unit of work?
70. What is the biggest threat to efficient performance of a distributed database system?
71. What is the difference between replication and propagation?
72. What is the difference between a data warehouse and a data mart?
73. Is it necessary to develop a backup and recovery plan for the data warehouse, since it is a read-only environment?
74. Describe the star schema and why it is an effective design for data warehouse databases.
75. What is the biggest DBA challenge when managing an Internet-enabled database?
76. What is the difference between the four types of JDBC database drivers?
77. What is metadata, and why is it important?
78. Name the three categories of performance tools, and briefly summarize the typical capabilities they provide.
79. Describe the difference between system time and business time in a temporal database system.
80. If you are asked by management to develop a cost justification for DBA tools, how can you minimize the amount of work it takes to produce such a justification?

A. Database Fundamentals

Most of the readers of this book should understand the basic concepts and fundamentals of database technology. However, many folks who think they understand the basics often do not have the knowledge and understanding they believe they have. Therefore, this appendix serves as a very brief introduction to the fundamentals of database management systems.

What Is a Database?

What is a database? The answer to this question may surprise some readers. Oracle is not a database; neither is DB2 or SQL Server. Each of these is a DBMS, or *database management system*. You can use Oracle, DB2, or SQL Server to create a database, but none of these themselves are databases. Many people, even skilled professionals, confuse the overall system—the DBMS—with the creation of the system—databases.

So, what is a database? A database is a structured set of persistent data. A phonebook is a database. However, within the world of IT, a database usually is associated with software. A simple database might be a single file containing many records, each of which contains the same set of fields where each field is a certain data type and length. In short, a database is an organized store of data where the data is accessible by named data elements.

A database is an organized store of data wherein the data is accessible by named data elements.

A DBMS is a software package designed to create, store, and manage databases. The DBMS software enables end users or application programmers to share data. It provides a systematic method of creating, updating, retrieving, and storing information in a database. DBMS products are usually responsible for data integrity, data access control, automated rollback, restart, and recovery.

You might think of a database as a file folder, and a DBMS as the file cabinet holding the labeled folders. You implement and access database instances using the capabilities of the DBMS. Your payroll application uses the payroll database, which may be implemented using a DBMS such as

IMS, DB2, Oracle Database 12g, or SQL Server.

Why is this distinction important? If we do not use precise terms in the workplace confusion can result, leading to over-budget projects, improperly developed systems, and lost productivity. Therefore, precision is important to us.

Why Use a DBMS?

The main advantage of using a DBMS is to impose a logical, structured organization on the data. A DBMS delivers economy of scale for processing large amounts of data because it is optimized for such operations.

A DBMS can be distinguished by the model of data upon which it is based. A *data model* is a collection of concepts used to describe data. A data model has two fundamental components: its *structure*, which is the way data is stored, and its *operations*, which is the way that data can be manipulated. The major DBMS products use four different data models:

1. Network (or CODASYL)
2. Hierarchical
3. Relational
4. Object oriented

There are four DBMS data models: hierarchical, network, relational, and object oriented.

The *network data model* is structured as a collection of record types and the relationships between these record types. All relationships are explicitly specified and stored as part of the structure of the DBMS. Another common name for the network model is CODASYL. CODASYL is named after the Conference on Data Systems Languages, the committee that formulated the model in the early 1970s. Data is manipulated using the location of a given record and following links to related records. IDMS is an example of a DBMS based on the network model.

The *hierarchical data model* arranges data into structural trees that store data at lower levels subordinate to data stored at higher levels. A hierarchical data model is based on the network model with the additional restriction that access to a record can be accomplished in only one way. IMS is an example of a DBMS based on the hierarchical model.

The *relational data model* consists of a collection of tables (more properly, *relations*) wherein the columns define the relationship between tables. The relational model is based on the mathematics of set theory. Contrary to popular belief, the relational model is not named after “relationships,” but after the relations of set theory. A *relation* is a set with no duplicate values. Data can be manipulated in many ways, but the most common way is through SQL. DB2, Oracle, and SQL Server are examples of DBMS products based on the relational model.

The *object-oriented (OO) data model* consists of a collection of entities, or objects, where each object includes the actions that can take place on that object. In other words, an object encapsulates data and process. With OO systems, data is typically manipulated using an OO programming language. Progress Software’s ObjectStore and Intersystems’ Cache are examples of DBMS products based on the OO model.

Each of these four data models is referred to as a data model for the sake of simplicity. In reality, only the relational and network models have any true, formal data model specification. Different models of data lead to different logical and structural data organizations. The relational model is the most popular data model because it is the most abstract and easiest to apply to data, while providing powerful data manipulation and access capabilities.

Most commercial DBMS implementations today are relational.

Advantages of Using a DBMS

Additionally, a DBMS provides a central store of data that can be accessed by multiple users, from multiple locations. Data can be shared among multiple applications, rather than having to be propagated and stored in new files for every new application. Central storage and management of data within the DBMS provides

Other Types of DBMS

Although the four data models discussed heretofore are the predominant types of DBMS, there are other types of DBMS with varying degrees of commercial acceptance.

A column-oriented DBMS, sometimes called a column store, is a

DBMS that stores its content by column rather than by row. This has advantages for data warehouses where aggregates are computed over large numbers of data items. Of course, a column-oriented DBMS is not based on any formal data model and can be thought of as a special physical implementation of a relational DBMS. Sybase IQ and Greenplum are examples of column stores. The NoSQL database system is another type of DBMS that is gaining market acceptance, usually in Big Data applications. Most NoSQL DBMS products are implemented as key-value data stores, where a piece of data is associated with a key. The data is not rigidly structured and does not have to conform to a schema as in a typical database design. NoSQL database systems are popular with organizations that face different data challenges from those that can be solved using traditional RDBMS solutions. Cassandra and MongoDB are examples of NoSQL key-value database systems.

There are other DBMS implementations as well, such as the inverted list structure of the original Adabas and even the dBase format popularized by the PC DBMSs, dBase II and dBase III.

- Data abstraction and independence
- Data security
- A locking mechanism for concurrent access
- An efficient handler to balance the needs of multiple applications using the same data
- The ability to swiftly recover from crashes and errors
- Robust data integrity capabilities
- Simple access using a standard API
- Uniform administration procedures for data

Levels of Data Abstraction

A DBMS can provide many views of a single database schema. A view defines what data the user sees and how that user sees the data. The DBMS provides a level of abstraction between the conceptual schema that defines the logical structure of the database and the physical schema that describes the files, indexes, and other physical mechanisms used by the database. Users function at the conceptual level—by querying columns within rows of tables,

for example—instead of having to navigate through the many different types of physical structures that store the data.

A DBMS makes it much easier to modify applications when business requirements change. New categories of data can be added to the database without disruption to the existing system.

Data Independence

A DBMS provides a layer of independence between the data and the applications that use the data. In other words, applications are insulated from how data is structured and stored. The DBMS provides two types of data independence:

- *Logical data independence*—protection from changes to the logical structure of data
- *Physical data independence*—protection from changes to the physical structure of data

As long as the program uses the API (application programming interface) to the database as provided by the DBMS, developers can avoid changing programs because of database changes.

SQL Is an API

The primary API to relational databases is SQL. In general, most application SQL statements need not change when database structures change (e.g., a new column is added to a table).

Data Security

Data security prevents unauthorized users from viewing or updating the database. The DBMS uses IDs and passwords to control which users are allowed access to which portions of the database. For example, consider an employee database containing all data about individual employees. Using the DBMS security mechanisms, payroll personnel can be authorized to view payroll data, whereas managers could be permitted to view only data related to project history.

Concurrency Control

A DBMS can serve data to multiple, concurrently executing user programs. This requires a locking mechanism to deliver concurrency control because the

actions of different programs running at the same time could conceivably cause data inconsistency. For example, multiple bank ATM users might be able to withdraw \$100 each from a checking account containing only \$150. A DBMS ensures that such problems are avoided because the locking mechanism isolates transactions competing for the same exact data.

Database Logging

The DBMS uses database logging to record “before” and “after” images of database objects as they are modified. It is important to note that the database log captures information about every data modification (except in circumstances as determined by the DBA). The information on the database logs can be used to undo and redo transactions. Database logging is handled transparently by the DBMS—that is, it is done automatically.

Ensuring Atomicity and Durability

A DBMS can be used to assure the all-or-nothing quality of transactions. This is referred to as [*atomicity*](#), and it means that data integrity is maintained even if the system crashes in the middle of a transaction. Furthermore, a DBMS provides recoverability. After a system failure, data can be recovered to a state that existed either immediately before the crash or at some other requisite point in time.

Data Integrity

The DBMS provides mechanisms for defining rules that govern the type of data that can be stored in specific fields or columns. Only data that conforms to the business rules will ever be stored in the database. Furthermore, the DBMS can be set up to manage relationships between different types of data and to ensure that changes to related data elements are accurately implemented.

Data Access

A DBMS provides a standard query language to enable users to interactively interrogate the database and analyze its data. For relational databases, this standard API is SQL, or *Structured Query Language*. However, SQL is not a requirement for a DBMS to be relational. Furthermore, many DBMS products ship with analytical tools and report writers to further simplify data access.

Summary

This section on DBMS fundamentals is necessarily brief because the focus of this book is on database administration, and most readers will find this material to be familiar. If you require additional details on the basic operations and qualities of DBMSs and databases, please refer to the Bibliography for an extensive list of DBMS-related books. My favorites include

- C. J. Date's *An Introduction to Database Systems*, 8th ed., for an academic and theoretical approach to the material
- Fabian Pascal's two books, *Practical Issues in Database Management* and *Understanding Relational Databases*, for an opinionated but informed approach to the topics
- Joe Celko's *Data & Databases: Concepts in Practice*, for a good practical overview of the topic
- The latest editions of Pratt and Adamski's *The Concepts of Database Management*, or Rob and Coronel's *Database Systems: Design, Implementation, and Management*, both of which provide a good high-level overview of DBMS concepts

A primary benefit of a DBMS is its ability to maintain and query large amounts of data while assuring data integrity and consistency. It offers transparent recovery from failures, concurrent access, and data independence. In fact, most modern computer applications rely on DBMS and database technology to manage data. Understanding the topic is of benefit to all IT professionals.

B. The DBMS Vendors

There are many DBMS vendors from which to choose. However, there are definite tiers in terms of popularity, support, and leadership in the DBMS market. In general, the marketplace can be broken down into the following groups:

- *The Big Three*—the three market leaders that constitute the greater part of the DBMS installed base as well as the bulk of any new sales
- *The second tier*—large DBMS vendors with stable products, but lagging behind the Big Three in terms of functionality and number of users
- *Other significant players*—other DBMS vendors with viable, enterprise-capable products
- *Open source*—DBMS products supported as open-source software (as opposed to by a single vendor)
- *Nonrelational*—vendors that supply prerelational DBMS products to support legacy applications
- *NoSQL*—nonrelational DBMS products for Big Data that are highly scalable to support modern Web applications
- *Object oriented*—vendors of ODBMS products that are used in conjunction with OO languages and development projects
- *PC based*—although many of the other vendors create PC versions of their DBMS products, these companies or products focus exclusively on the PC platform

The following sections present a bit more information on each of the marketplace groups.

The Big Three

Oracle Corporation, IBM Corporation, and Microsoft are the leading DBMS vendors, in that order. The installed base for Oracle is greater than that of any other DBMS, and Oracle is probably used by more people than any other DBMS. IBM's DB2 is in second place and picking up steam. With DB2's near monopoly on the mainframe and its growing UNIX and Windows

installed base, IBM can claim to be used by almost as many users as Oracle. In third place overall, but in first place on Windows machines, is Microsoft, with Microsoft SQL Server. Microsoft SQL Server runs only on Windows platforms.

Contact information:

Oracle Corporation
500 Oracle Parkway
Redwood Shores, CA 94065

www.oracle.com

IBM Corporation
New Orchard Road
Armonk, NY 10504

www.ibm.com

Microsoft
1 Microsoft Way
Redmond, WA 98052-6399

www.microsoft.com

The Second Tier

The second-tier DBMS providers category comprises companies that were market leaders in the 1980s. Sybase was a leading innovator in database technology, introducing such concepts as the client/server paradigm, stored procedures, and more. Indeed, Microsoft SQL Server is based on Sybase technology. Sybase was acquired by German packaged application vendor SAP in 2010.

Teradata is another innovator in database technology that focuses its business model primarily on data warehousing and analytics. Teradata was incorporated in 1979 and was formerly a division of NCR Corporation, but it became an independent entity with a spinoff from NCR in October 2007.

Informix is another former leader that fell on hard times with a misguided development strategy. All of Informix's DBMS assets were acquired by IBM Corporation in the summer of 2001, and Informix continues to be supported by IBM and used by many performance-minded organizations.

Contact information:

Sybase Inc. (an SAP Company)

6475 Christie Avenue
Emeryville, CA 94608
www.sybase.com

Teradata Corporation
10000 Innovation Drive
Dayton, OH 45342
www.teradata.com

Other Significant Players

Besides the vendors listed previously, several additional vendors offer relational products with a high degree of functionality and with a reasonable number of installed users. Foremost among these is Ingres, currently marketed by Actian Corporation. Ingres began its life as the first relational DBMS implementation at the University of California, Berkeley. The Ingres code was enhanced into a commercial product by Relational Technologies. Relational Technologies became Ingres Corporation, which was purchased by ASK Corporation, which was then purchased by Computer Associates (now CA Technologies). Computer Associates eventually spun off Ingres as a separate company, which open-sourced Ingres. That company renamed itself Actian in 2011. Ingres still boasts a large number of users, particularly in Europe.

Another significant player is Software AG's Adabas and Tamino DBMS products. Adabas has its roots as an inverted-list DBMS that was augmented to become relational. Tamino is Software AG's new XML-based DBMS.

Contact information:

Actian Corporation
500 Arguello Street, Suite 200
Redwood City, CA 94063
www.actian.com

Software AG
Uhlandstr. 12
64297 Darmstadt
Germany
www.softwareag.com

Open-Source DBMS Offerings

The rapid acceptance and media glare surrounding Linux has enlivened the open-source community. The term *open source* refers to software that users are free to run, copy, distribute, study, change, and improve. Often “open source” is interpreted to mean free software. This is understandable, but the open-source concept of free is closer to “liberty” than to “no charge.” Open-source software adheres to the following beliefs:

- Users are free to run the program, for any purpose.
- Users are free to inspect the actual source code of the program to determine how it works.
- Users are free to modify and adapt the software to their specific needs.
- Users are free to distribute copies to whomever.
- Users are free to release their code improvements to the public, so that the whole community benefits.

There are two leading open-source DBMS products: PostgreSQL and MySQL. In addition to becoming a commercial product, Ingres was enhanced at Berkeley, eventually turning into a second-generation DBMS named Postgres, and eventually a SQL version named PostgreSQL. There are several PostgreSQL distributors, or it can be downloaded for free at www.postgresql.org/.

MySQL is a true multiuser, multithreaded SQL database server. It is available as a free download from www.mysql.com, but it is not technically an open-source DBMS. Because you do not get access to the source code, you cannot modify it. Additionally, you must pay for a license if you sell MySQL, if you charge for administering a server, or if you include MySQL as a component of another product.

Oracle Corporation acquired MySQL in 2009 but continues to support and market it much as MySQL was before the acquisition.

Nonrelational DBMS Vendors

Prior to the commercial acceptance of relational technology in the early 1980s, the DBMS market was primarily mainframe based and was dominated by two players: IBM and Cullinet. IBM’s IMS is a hierarchic DBMS that organizes data into tree structures. Cullinet marketed IDMS, a CODASYL DBMS that organized data into network structures. Cullinet was acquired by Computer Associates in the late 1980s.

NoSQL DBMS Vendors

Another, more recent form of nonrelational DBMS is the NoSQL database system. At a high level, NoSQL implies nonrelational, distributed, flexible, and scalable. Many are also open source. NoSQL grew out of the perceived need for “modern” database systems to support Web initiatives. Additionally, some common attributes of NoSQL DBMSs include the lack of a schema, simplicity of use, replication support, and an “eventually consistent” capability (instead of the typical ACID transaction capability). It really does not mean no SQL support. Indeed, today the “no” is more commonly defined as “not only.”

There are a large number of NoSQL “database systems” available for use, but the more popular ones include

- *CouchDB*—a document-oriented database that can be queried and indexed in a MapReduce fashion using JavaScript. CouchDB also offers incremental replication with bidirectional conflict detection and resolution.
- *MongoDB*—a scalable, high-performance, open-source, document-oriented database system.
- *Hadoop and HBase*—The Apache Hadoop project develops open-source software for reliable, scalable, distributed computing. HBase is the Hadoop database system. It supports random, real-time read/write access to “Big Data.”

The Web site at <http://nosql-database.org/> offers a nice rundown of the commercially available NoSQL database systems available.

Object-Oriented DBMS Vendors

At one point in the late 1980s and early 1990s, it looked as if object-oriented database management systems (ODBMS) were going to gain rapid acceptance. For many reasons (difficult to query, no solid data model, etc.), this technology was not widely adopted. Indeed, many of the ODBMS vendors have repositioned their products as e-business solutions or database components for object-oriented application developers. However, there are niche pockets of ODBMS users.

The following vendors still offer ODBMS products and components:

- Object Store from Progress Software at www.progress.com
- Ontos at www.ontos.com
- Poet at www.poet.com

PC-Based DBMS Vendors

Simpler to use than full-blown enterprise DBMS products, PC-based DBMS products provide structured storage for PC-based data and applications. The following vendors and products are the market leaders in the PC-based DBMS market:

- dBase at www.dbase.com
- FileMaker at www.filemaker.com
- Lotus Approach at www.ibm.com/software/lotus/products/smartsuite/approach.html (included in Lotus SmartSuite)
- Microsoft Access at www.microsoft.com/office/access/
- Paradox at www.corel.com (included in the Professional edition of WordPerfect Office)

C. DBA Tool Vendors

In [Chapter 23](#) we covered the wide variety of tools available to ease the burden of database management and administration. Many of the world's largest independent software vendors (ISVs) supply DBA tools. This appendix lists the major players.

The Major Vendors

Although there are many ISVs that manufacture and market DBA tools and utilities, the four vendors that follow are the leading independent DBA tool vendors. These vendors distinguish themselves because they offer tools that work for more than a single DBMS and in more than one tool category. When evaluating DBA tools for a heterogeneous shop with multiple DBMSs, these are the first vendors you should consider.

BMC Software
2101 City West Blvd.
Houston, TX 77042
www.bmc.com

CA Technologies, Inc.
One CA Plaza
Islandia, NY 11749
www.ca.com

Embarcadero Technologies
100 California Street, 12th Floor
San Francisco, CA 94111
www.embarcadero.com/

Of course, all of the DBMS vendors themselves also market add-on DBA tools for their respective DBMS products. Refer to [Appendix B](#) for a listing of the major DBMS vendors and contact information for each.

Other DBA Tool Vendors

In addition to the major DBA tool vendors mentioned in the preceding section, each of the following vendors markets add-on DBA tools. In contrast to the vendors listed previously, these vendors are likely to specialize in

certain niche areas. For example, some may market only performance management tools, and others may offer tools that work only on Microsoft SQL Server.

The Web sites listed for these vendors provide additional information on the type of tools offered.

- Allen Systems Group: www.asg.com/
- Application Security: www.appsecinc.com/
- Aquafold: www.aquafold.com/
- Bradmark Technologies: www.bradmark.com/
- Bunker Hill Corporation: www.bunkerhill.com/
- Candle Corporation: www.candle.com
- CDB Software: www.cdbsoftware.com
- Cogito: www.cogito.co.uk/
- Compuware Corporation: www.compuware.com
- Confio Software: www.confio.com/
- DataBee: www.databee.com/
- DBE Software: www.dbesoftware.com/
- DBI Software: www.dbisoftware.com/index.php
- dbMaestro: www.dbmaestro.com/
- Hit Software: www.hit.com
- Idera: www.idera.com/Content/Home.aspx
- InfoTel Corporation: www.infotelcorp.com
- NetIQ: www.netiq.com
- Quest Software: www.quest.com (purchased by Dell in 2012)
- Red Gate Software: www.red-gate.com
- Relational Architects: www.relarc.com
- Responsive Systems: www.responsivesystems.com
- Rocket Software: www.rocketsoft.com
- SEGUS Inc.: www.segus.com
- Softbase Systems Inc.: www.softbase.com
- SoftwareOnZ LLC: www.softwareonz.com

- White Sands Technology, Inc.: www.whitesands.com

These are not the only vendors marketing add-on DBA tools, but they are some of the bigger and more successful ISVs. Additional tool categories and related ISVs are listed in the following sections.

Data Modeling Tool Vendors

Computer Associates International

www.ca.com

Product: ERwin

Datanamic

www.datanamic.com

Product: DeZign for Databases

Embarcadero Technologies

www.embarcadero.com

Product: ER/Studio

Grandite

www.silverrun.com

Product: SILVERRUN

IBM Corporation

www.ibm.com

Product(s): Rational System Architect, InfoSphere Data Architect

Microsoft Corporation

www.microsoft.com

Product: Visio Enterprise

ModelRight

www.modelright.com

Product: ModelRight

Oracle Corporation

www.oracle.com

Product: Oracle Designer

Powersoft (Open Tools Division of Sybase)

www.sybase.com/products/modelingdevelopment/powerdesigner

Product: Power Designer

Visible Systems

www.visible.com/

Product: Visible Analyst, Visible Advantage

Repository Vendors

Adaptive Ltd.

www.adaptive.com/

Product: Adaptive Information Manager

Allen Systems Group

www.asg.com

Product: ASG-Rochade, ASG-Manager, ASG-Vista

CA Technologies, Inc.

www.ca.com

Product: CA Repository

IBM Corporation

www.ibm.com

Product: ClearCase

Data Movement and Business Intelligence Vendors

The following companies offer data movement and ETL tools:

- CoSort/IRI: www.cosort.com/
- Informatica: www.informatica.com
- Information Builders: www.informationbuilders.com
- Serena: www.serena.com/
- SQData: www.sqdata.com
- Treehouse Software: www.treehouse.com
- Vision Solutions: www.visionsolutions.com/

The following companies offer business intelligence and OLAP tools:

- Actuate: www.actuate.com
- Cardett Associates: www.cardett.co.nz/
- Computer Associates: www.ca.com
- IBM Corporation (Cognos): www.ibm.com
- Oracle Corporation (Hyperion): www.oracle.com

- SAP (Business Objects): www.sap.com, www.sap.de
- SAS Institute: www.sas.com
- Visual Insights: www.visualinsights.com/

Once again, all of the major DBMS vendors offer add-on ETL and business intelligence tools for their respective DBMS products (and frequently for heterogeneous ETL and BI, too).

D. DBA Web Resources

DBAs are not alone in the Internet-connected world. DBAs have many online resources at their disposal on the Web. By virtue of being Internet connected, DBAs have access to the vast knowledge and experience of their peers. However, to take advantage of these online resources, the DBA has to know that the resources exist, how to gain access to them, and where to find them. This appendix will discuss several of the Internet resources available to DBAs. DBAs who do not take advantage of the Internet resources at their disposal are doing themselves a disservice.

Usenet Newsgroups

When discussing the Internet, many folks limit themselves to the World Wide Web. However, there are many components to the Internet. One often-overlooked resource is the Usenet newsgroup, which can be a very fertile source of expert information. Usenet, an abbreviation for User Network, is a large collection of discussion groups, called newsgroups. Each newsgroup is a collection of postings pertaining to a single, predetermined topic. Newsgroup names usually reflect their focus. For example, `comp.databases.ibm-db2` contains discussions about the DB2 family of products.

Using newsreader software, any Internet user can access a newsgroup. There are a number of options for newsreader software that can be downloaded and used free of charge. Consider reviewing the options available at www.newsgroupreviews.com/.

There are many newsgroups that focus discussion on database and database-related issues. [Table D.1](#) shows some of the most pertinent newsgroups of interest to the DBA.

Table D.1. Database-Related Usenet Newsgroups of Interest to DBAs

Newsgroup Name	Description
comp.client-server	Information on client/server technology
comp.compression.research	Information on research in data compression techniques
comp.data.administration	Discussion of data modeling and data administration issues
comp.databases	Issues regarding databases and data management
comp.databases.ibm-db2	Information on IBM's DB2 family of products
comp.databases.informix	Information on the Informix DBMS
comp.databases.ms-sqlserver	Information on Microsoft's SQL Server DBMS
comp.databases.object	Information on object-oriented database systems
comp.databases.olap	Information on data warehouse online analytical processing
comp.databases.oracle.marketplace	Information on the Oracle market
comp.databases.oracle.server	Information on the Oracle RDBMS
comp.databases.oracle.tools	Information about add-on tools for Oracle
comp.databases.oracle.misc	Miscellaneous Oracle discussions
comp.databases.sybase	Information on the Sybase Adaptive Server RDBMS
comp.databases.theory	Discussions on database technology and theory
comp.unix.admin	UNIX administration discussions

Of course, thousands of other newsgroups exist. You can use your newsreader software to investigate the newsgroups available to you and to gauge the quality of their discussions.

Mailing Lists

Another useful Internet resource for DBAs is the mailing list. Mailing lists are a kind of community bulletin board. You can think of mailing lists as somewhat equivalent to a mass mailing. However, mailing lists are not spam because users must specifically request to participate before they will receive any mail. This is known as *opting in*.

There are numerous mailing lists available on the Internet, and they operate by use of a list server. A *list server* is a program that automates the mailing list subscription requests and messages. The two most common list servers are Listserv and Majordomo. Listserv is also a common synonym for mailing list, but it is actually the name of a particular list server program.

Simply by subscribing to a mailing list, you can get information sent directly to your e-mail in-box from the list server. The information that you will receive varies from news releases, to announcements, to questions, to answers. This information is very similar to the information contained in a newsgroup forum, except that it comes directly to you via e-mail. Users can

also easily respond to mailing list messages, enabling communication with every subscribed user. Responses are sent back to the list server as e-mail, and the list server sends the response out to all other members of the mailing list.

To subscribe to a mailing list, simply send an e-mail to the appropriate address requesting a subscription. Several Web sites, including the following, catalog the available Internet mailing lists:

- www.lsoft.com/lists/listref.html
- www.listtool.com

Of course, none of these sites track every single mailing list available to you. Vendors, consultants, Web portals, and user groups also support mailing lists of various types. The only way to be sure you know about all the useful mailing lists out there is to become an actively engaged member of the online community. [Table D.2](#) provides details on a few popular database-related mailing lists for DBAs.

Table D.2. Database-Related Mailing Lists of Interest to DBAs

Mailing List Name	Subscription Address	Description
ORACLESIG	E-mail <code>LISTSERV@LISTS.BLACKBOARD.COM</code> with the command <code>SUBSCRIBE ORACLESIG</code>	Oracle DBA special-interest group
DB2-L	E-mail <code>LISTSERV@WWW.IDUG.ORG</code> with the command <code>SUBSCRIBE DB2-L</code>	Discussion about the DB2 family of products
SQL2K8	E-mail <code>lists@sswug.org</code> with the command <code>SUBSCRIBE SQL2K8</code>	Discussion of Microsoft SQL Server and related products

Although mailing lists are not as popular as they once were, some are still very active and vital resources for database information.

Web Sites, Blogs, and Portals

Of course, the Web is also a very fertile source of database- and DBA-related information. However, tracking things down on the Web can sometimes be difficult—especially if you do not know where to look. Several good sources of DBMS information can be found by reviewing the Web sites of DBMS vendors, DBA tool vendors, magazines, and consultants

Vendor Web Sites

One of the best DBA resources is the vendor Web site for the DBMS software you use. On such sites you can find up-to-date news, information on release schedules and bug fixes, downloadable manuals, and technical support. Be sure to bookmark the site for your DBMS vendor and visit regularly to keep tabs on what's going on with your favorite database management systems.

The Web sites for the major DBMS vendors are

- IBM DB2 at www.ibm.com/software/data/db2/
- Informix at www.ibm.com/software/data/informix/
- Microsoft SQL Server at www.microsoft.com/sql/default.asp
- Oracle at www.oracle.com/us/products/database/index.html
- Sybase at www.sybase.com

Another useful resource for DBAs can be found on DBA tool vendor Web sites. DBA tool vendors employ many DBMS experts and former DBAs who are willing to share tips, techniques, and scripts in order to promote their company and its products. [Appendix C](#) lists several of the most popular third-party tool vendors along with their Web address and contact information.

Magazine Web Sites

Many useful database and DBA-related articles are published every week in technical magazines of every sort. Several magazines focus almost exclusively on database- and DBA-related issues. Even better, many of these publications are provided free of charge, in return for filling out a subscription request with company and demographic information.

Additionally, most print magazines also host Web sites that offer many, if not all, of their articles online. Some of this content is available to subscribers only, but many periodicals offer some content to all browsers. Some of the better magazine Web sites out there include

- *Oracle Magazine*: www.oracle.com/oramag/index.html
- *IBM Data Management* magazine: www.ibm.com/developerworks/data/dmmag/
- *Database Trends and Applications*: www.dbta.com
- *Information Management*: www.information-management.com/
- *SQL Server Pro* magazine: www.sqlmag.com

- *Teradata Magazine*: www.teradatamagazine.com/

Consultant Web Sites

There is a boatload of consultants who specialize in DBMS technology. However, a few consultants have achieved guru status because they are willing to share their knowledge and experience at conferences and events, and by writing articles. Many of these premier consultants have set up content-rich Web sites containing numerous documents, presentations, and tips for anyone visiting the site. The consultants do this to promote their business and gain leads for future business.

Some of the leading consultant Web sites containing interesting and useful information for DBAs include

- Yevich, Lawson & Associates, focusing primarily on DB2: www.ylassoc.com
- TUSC, focusing on Oracle: www.tusc.com
- DBA Direct, focusing on database administration: www.dbadirect.com
- Fabian Pascal's site, devoted to debunking erroneous beliefs about relational and other database technologies: www.dbdebunk.com
- The Pythian Group, focusing on DBA and consulting services: www.pythian.com/

And, of course, you can check out the author's Web site at www.craigsmullins.com for an exhaustive online catalog of his published writing and more.

Blogs

Another fertile source of database information on the Web is the blog. A blog is a particular type of Web site (or a portion of a Web site) that is frequently updated and less formal than the typical technology Web site or article.

The term *blog* is a blending of the phrase *Web log*. Blogs typically are maintained by a single individual with regular entries of commentary, descriptions of events, and other material. Entries, known as blog postings, usually are displayed in reverse chronological order. Most blogs are also interactive, containing sections where visitors can leave comments. For some blogs, the comments section can be more active than the content area. The more comments, the more fun and informative the blog usually will be.

According to Wikipedia, as of February 2011, there were over 156 million blogs in existence. Some of the more entertaining and informative blogs that cover database- and DBA-related material include

- DB2 Portal Blog: <http://db2portal.blogspot.com>
- Getting the Most out of DB2 for z/OS: <http://it.toolbox.com/blogs/db2zos>
- An Expert's Guide to DB2 Technology: <http://it.toolbox.com/blogs/db2luw>
- Oracle Scratchpad: <http://jonathanlewis.wordpress.com/2011/12/29/i-wish-4/>
- An Expert's Guide to Oracle Technology: <http://it.toolbox.com/blogs/oracle-guide>
- Richard Foote's Oracle Blog: <http://richardfoote.wordpress.com/>
- SQL Is Your Friend: <http://it.toolbox.com/blogs/sqlisyourfriend>
- SQL Marklar: www.networkworld.com/community/mccown
- SQL Rockstar: <http://thomaslarock.com/>
- Data and Technology Today: <http://datatechnologytoday.wordpress.com/>
- Obsessive-Compulsive Data Quality: www.ocdqblog.com/
- Data Governance: <http://datagovernanceblog.com/>

Seek out blog aggregation Web sites for help in finding and keeping up-to-date on specific blogs. For example, the best source for finding DB2 blogs is Planet DB2, a blog aggregator service at www.planetdb2.com.

And finally, consider bookmarking Pythian's Log Buffer, a Carnival of the Vanities blog at www.pythian.com/news/. A carnival blog can be thought of as a blog of blogs. Log Buffer, in particular, offers an interesting way to keep up-to-date on what is being said in the blogosphere regarding database management. It can be useful if you want to quickly track a whole bunch of different database blogs. According to the Pythian Group, Log Buffer “. . . is the place DBAs and others interested in databases come to read and write the news of what their colleagues around the world are working on: how they're solving problems, using their skills, and what they're thinking about.”

Database Portals

There are many Web portals that offer a convenient collection of database- and DBA-related information. A portal typically provides a vast amount of information on a single topic, as well as links to related information, a search engine, customization options, and the ability for users to be notified by e-mail when relevant content is posted. The following Web sites function as portals for database professionals.

The Database Site (www.thedatabasesite.com) is a source of useful information about database management, administration, and development issues and solutions. The site is a combination portal and Web-based magazine with regular new content and links to vital database-related information and sites.

TechTarget publishes SearchOracle (www.searchoracle.com) and SearchSQLServer (<http://searchsqlserver.techtarget.com>), as well as their sister site SearchDataManagement (<http://searchdatamanagement.techtarget.com>). These sites are portals containing loads of useful information, including tips, articles, discussion boards, “ask the expert” sections, and more.

The SQL Server Worldwide User Group (www.sswug.org/) is another very helpful Web portal for DBAs. The name seems to indicate that the portal specializes in SQL Server, and that is the predominant subject area of SSWUG. However, the portal also covers a very broad scope of database-related subjects, including DB2, Oracle, XML, security, and others.

The Data Administration Newsletter (www.tdan.com) is another combination portal/magazine site, which focuses on data administration and data architecture issues.

There are numerous other portals that provide useful information for DBAs. Some of the best focus on a single DBMS technology, whereas others offer heterogeneous information.

Other Web Sites

Many other Web sites may prove useful or interesting to DBAs. One of the best things about the Internet is its ability to provide a forum for anyone with information to share. New and intriguing sites are popping up every day. Let’s take a look at some of the best Web sites out there for DBAs.

Providing a very useful service for DBAs is DBAjobs.com (www.DBAjobs.com). This site’s mission is “to advance the careers of all

database professionals and to recruit the best database talent for its clients.” Be sure to visit this site for access to many database-related jobs as well as for assistance in preparing your résumé, interviewing tips, salary guidelines, and database news.

Sites hosting information on industry standards can be quite helpful, too. For example, ANSI (American National Standards Institute), ISO (International Standards Organization), NIST (National Institute of Standards and Technology), and SEI (Software Engineering Institute) all provide Web sites:

- ANSI: www.ansi.org
- ISO: www.iso.ch
- NIST: www.nist.gov
- SEI: www.sei.cmu.edu/

Perhaps closer to the technical information required by a DBA is the SQL standard Web site at www.jcc.com/sql.htm. This site purports to be a central source of information about SQL standards. It contains links to various other sites that document the SQL standard.

DAMA, the Data Management Association, is a coalition of professionals dedicated to the advancement of data resource management. The DAMA Web site (www.dama.org) documents the conferences and meetings sponsored by DAMA, and also offers links, resources, newsletters, and publications on data resource management.

User groups have been founded for each of the major DBMS vendors. These groups sponsor annual international conferences and promote the sharing of information by database professionals. The Web sites for these groups offer useful information on conference locations and dates, as well as articles, presentations, and tips:

- International DB2 User Group: www.idug.org
- International Oracle Users Group: www.ioug.org/
- Professional Association for SQL Server: www.sqlpass.org
- International Informix Users Group: www.iiug.org/
- International Sybase User Group: www.isug.com/

There are literally thousands of other Web sites that might be useful to you as you embark on your career as a DBA. Here are a few of the ones I have

found to be worthwhile:

- Agile Modling: www.agilemodeling.com/
- Business Intelligence Network: www.b-eye-network.com/
- The Data Warehousing Information Center: www.dwinfocenter.org/
- The Data Warehousing Institute: www.tdwi.org/
- Object Management Group: www.omg.org
- The OLAP Council: www.olapcouncil.org/
- The Open Group: www3.opengroup.org/
- Storage Networking Industry Association: www.snia.org
- The Transaction Processing Council: www.tpc.org/
- The XML portal: www.xml.com
- Web Farming site (data warehouse): www.webfarming.com

All of these sites are very useful for obtaining up-to-date information about DBMS releases and versions, management tool offerings, and the like, but bear in mind that the information on this type of site can be very biased. For information that is more likely to be unbiased, you should favor the Web portals over the vendor and consultant sites.

E. Sample DBA Job Posting

This appendix provides an example job posting for a database administrator. It can be used as a template for organizations looking to post a DBA job opening or by DBAs interested in the types of job responsibilities organizations look for when hiring.

Job Posting

Database Administrator (DBA)

Assists in the development of data models (perhaps in conjunction with a data architect, perhaps solo) for new projects within IT. Has primary responsibility for the definition, creation, and maintenance of enterprise databases and database structures within the organizations. Analyzes, tests, and implements physical database design supporting various business applications (including base definition, structure, documentation, long-range requirements, operational guidelines). Ensures data recovery, maintenance, data integrity, and space requirements for physical databases are met through formulations and monitoring of policies, procedures, and standards relating to database management. Provides a high degree of coordination among user, technical, application, and corporate entities to ensure effective data management as a corporate enterprise. Above all, ensures solutions meet business objectives and establishes and maintains a high level of user trust and confidence in IT's knowledge of and concern for users' business needs. Has a willingness to work with customers of varying levels of technical expertise in a high-pressure, complex environment.

Responsibilities

- Provide first-level applications support: Resolve problems and answer questions related to the RDBMS and platform environment (e.g., UNIX, Windows, Linux, z/OS); contact vendor support when necessary; facilitate application development and client problem resolution with optimum speed and efficiency.
- RDBMS performance monitoring and tuning: Closely monitor performance; identify problems and implement solutions; ensure the database is running at optimum speed and efficiency.

- Maintain database backup/recovery environment: Assure that data recovery time objectives are established and met; verify implementation of correct backup strategy; develop recovery procedures; support recovery from loss of data caused by user or system error; respond quickly to refresh data.
- Install and upgrade RDBMS and related products: Keep RDBMS at current release and maintenance levels; install supporting products; provide current, stable production and development environments.
- Reorganize database structures as needed: Automate procedures at regular intervals; use tools to reorganize or defragment database tablespaces, tables, and indexes; provide an efficient database environment to maximize productivity and improve performance.
- Perform application refreshes on request: Refresh data from one environment to another at the request of the application development staff; execute necessary tasks to accomplish this, including export, drop, create, and import; build and schedule a script to accomplish refreshes to minimize impact; provide application development staff optimal development environment with current data.
- Guide and oversee the database design process as needed: Research and recommend optimal design criteria; publish standards and work with development staff; standardize environment and improve efficiency.
- Review and implement all changes to database structures, while ensuring the integrity of the data and maintaining data availability as required by service-level agreements.
- Capacity planning and reporting: Monitor file system space, database allocation, etc.; verify efficient use of disk storage; reclaim unused space; optimize space allocation and avoid out-of-space conditions.
- Create, develop, and review the physical model created from the logical model, ensuring the detailed physical design of data structures and databases complies with standards, guidelines, and procedures recommended for all aspects of data modeling, database design, and maintenance; ensure proper data modeling and design paradigms are applied.
- Provide the appropriate tools, procedures, and facilities for creating and maintaining the entity data enterprise model.

- Be capable of producing entity-relationship diagrams; data flow diagrams; database normalization schemata; logical-to-physical database mapping; DDL and database commands.
- Manage the implementation and use of the enterprise data dictionary.
- Lead design reviews for database designs, application code, and pre- and post-implementation capacity meetings; estimate sizes of new databases.
- Work closely with systems programmers and systems administrators to ensure that the database system is operating in accordance with required system usage; provide assistance and guidance in setting appropriate system parameters.
- Develop rules, procedures, and standards for the access and maintenance of shared data resources.
- Review overall physical database structures for data integrity, performance quality, recoverability, maintenance, and space requirement considerations.
- Communicate complex technical information in a concise and articulate manner.
- Participate in the evaluation and recommendation of DBA tools and new DBMS technologies.
- Serve as business liaison to database vendors.
- Be responsible for the integrity of production test and QA database environments.
- Perform impact analyses, research products, configure software, and establish standards and procedures that support the company's database environment needs.
- Be responsible for application-oriented database administration and technical support for the various development, training, testing, and production relational database environments; focus on application-specific database issues to ensure optimal application performance and integrity.
- Develop utilities for monitoring and evaluating data quality.
- Assist in the definition and development of database documentation and standards.

- Conduct application transaction volume and traffic analysis, and interpret impact on database and network performance.
- Development projects: Interact with analysts and developers, understand system requirements from conceptual models, analyze transactions and data volumes, and design the physical database.

Skills and Knowledge

- Thorough understanding of the relational database model and solid theoretical knowledge of relational database techniques; skilled in ability to see the big picture and conceptualize and document creative solutions
- Experience in designing, modeling, developing, and supporting large RDBMS systems
- Extensive experience in logical and physical database design, hands-on data modeling techniques, and development methodology; physical schema design of tablespaces, rollback segments, and data files
- Experience in the design of data extract, data migration, and data queuing facilities
- Proficiency in the concepts and practice of database normalization techniques
- Strong working knowledge of state-of-the-art database concepts/design/deployment, recovery techniques, performance monitoring and tuning, Structured Query Language, relational database design techniques, JCL
- Knowledge of server-side development issues surrounding triggers, stored procedures, and packages
- Strong customer focus and management of client expectations; ability to establish and maintain a high level of user trust and confidence in IT's knowledge of and concern for users' business needs
- Excellent oral and written communications skills
- Ability to present and explain technical information to diverse types of audiences (management, users, vendors, technical staff) in a way that establishes rapport, persuades others, and gains understanding
- Ability to write clear proposals and technical documents
- Strong interpersonal skills at all levels of management and ability to

motivate employees/teams to apply skills and techniques to solve dynamic problems; excellent teamwork skills

- Solid project management skills, effectively managing multiple small to large projects in a cross-functional environment
- Ability to weigh various suggested technical solutions against the original business needs and choose the most cost-effective solution
- Ability to work with long-established senior staff, provide guidance, and assign and follow up on tasks
- Proficiency in several RDBMS-specific middleware tools and utilities to support such activities as monitoring, database migration, DDL management, and replication
- Experience with cross-platform networking in a TCP/IP environment
- Extensive knowledge of SQL, including programming and ways to optimize SQL code
- Solid programming knowledge of at least one 3GL/4GL/OO4GL language (e.g., C/C++, Perl, UNIX shell, Visual Basic)
- Knowledge of Web-enabled GUIs is a plus (e.g., Oracle JDeveloper and WebSphere Studio)

Experience

- Bachelor's degree in computer science, engineering, or related discipline; master's degree is desirable; equivalent experience acceptable
- 3+ years in-depth experience in the design, administration, programming, and support of RDBMSs (e.g., Oracle, SQL Server, DB2) in one or more computing environments (e.g., UNIX, z/OS, Windows, Linux, Web)
- Ability to handle basic systems administration tasks to fulfill backup role if needed
- Demonstrated experience working well with customers of varying levels of technical expertise in high-pressure situations and complex environments
- Highly desirable that individual has earned certification as a developer or DBA, for example, Oracle DBA Certified Master (OCM), Oracle DBA Certified Professional (OCP), Microsoft Certified Database

Administrator (MCDBA), Certified MySQL Professional (MySQL-Prof), IBM DB2 Universal Database Certified Solutions Expert (DB2), etc.

Working Relationships

- Has close working relationship with systems team and peers in other areas of IT
- Works closely with senior systems programmer, RDBMS and systems administrators
- Has contact with application users, primarily in problem resolution
- Has frequent interaction with application development staff as facilitator and client problem resolver; works collaboratively on such activities as performance problems, hardware/software upgrades, application program design reviews, new application setup and implementation, application workload growth prediction and measurement, etc.
- Is key participant in high-level technical strategy meetings, major upgrade decisions, capacity planning data collections, performance reporting, new application designing, etc.
- Has close working relationships with peers at external organizations

Bibliography

The books listed in this bibliography can be used as guides to further education and research on the topics of database management, database administration, and other data-related technologies.

Database Management and Database Systems

Abiteboul, Serge, et al. *Foundations of Databases*. Reading, MA: Addison-Wesley (1995). ISBN 0-201-53771-0.

Ambler, Scott W. *Agile Database Techniques: Effective Strategies for the Agile Software Developer*. New York, NY: John Wiley & Sons (2003). ISBN 0-471-20283-5.

Ambler, Scott W., and Pramod J. Sadalage. *Refactoring Databases: Evolutionary Database Design*. Boston, MA: Addison-Wesley (2006). ISBN 0-321-29353-3.

Atre, Shaku. *Database: Structured Techniques for Design, Performance, and Management*. 2nd ed. New York, NY: John Wiley & Sons (1988). ISBN 0-471-85251-1.

Ault, Mike, et al. *Database Benchmarking: Practical Methods for Oracle and SQL Server*. Kittrell, NC: Rampant Techpress (2006). ISBN 0-9776715-3-4.

Bell, David, and Jane Grimson. *Distributed Database Systems*. Wokingham, England: Addison-Wesley (1992). ISBN 0-201-54400-8.

Blaha, Michael R. *A Manager's Guide to Database Technology*. Upper Saddle River, NJ: Prentice Hall (2001). ISBN 0-13-030418-2.

Bobak, Angelo R. *Distributed and Multi-Database Systems*. New York, NY: Bantam Intertext (1993). ISBN 0-553-09156-5.

Bontempo, Charles J., and Cynthia Maro Saracco. *Database Management Principles and Products*. Upper Saddle River, NJ: Prentice Hall (1995). ISBN 0-13-380189-6.

Bradley, James. *Introduction to Database Management in Business*. New York, NY: Holt, Rinehart, and Winston (1983). ISBN 0-03-061693-X.

- Brathwaite, Kenmore S. *The Data Base Environment: Concepts and Applications*. New York, NY: Van Nostrand Reinhold (1990). ISBN 0-442-00300-5.
- Brathwaite, Kenneth S. *Systems Design in a Database Environment*. New York, NY: McGraw-Hill (1989). ISBN 0-07-007250-7.
- Burleson, Donald K. *Managing Distributed Databases*. New York, NY: John Wiley & Sons (1994). ISBN 0-471-08623-1.
- Castano, Silvana, et al. *Database Security*. Wokingham, England: Addison-Wesley (1994). ISBN 0-201-59375-0.
- Celko, Joe. *Data & Databases: Concepts in Practice*. San Francisco, CA: Morgan Kaufmann (1999). ISBN 1-55860-432-4.
- Chisholm, Malcolm. *Managing Reference Data in Enterprise Databases*. San Francisco, CA: Morgan Kaufmann (2001). ISBN 1-55860-697-1.
- Chorafas, Dimitris. *Handbook of Database Management and Distributed Relational Databases*. Blue Ridge Summit, PA: Tab Books (1989). ISBN 0-8306-3253-0.
- Chorafas, Dimitris N., and Heinrich Steinmann. *Solutions for Networked Databases*. San Diego, CA: Academic Press (1993). ISBN 0-12-174060-9.
- Codd, E. F. *The Relational Model for Database Management Version 2*. Reading, MA: Addison-Wesley (1990). ISBN 0-201-14192-2.
- Connolly, Thomas, and Carolyn Begg. *Database Systems: A Practical Approach to Design, Implementation, and Management*. 4th ed. Harlow, England: Addison-Wesley (2004). ISBN 978-0-321-29401-2.
- . *Database Systems: A Practical Approach to Design, Implementation, and Management*. 2nd ed. Harlow, England: Addison-Wesley (1998). ISBN 0-201-34287-1.
- Courtney, James F., and David B. Paradise. *Database Systems for Management*. 2nd ed. Homewood, IL: Irwin (1992). ISBN 0-256-08229-4.
- Date, C. J. *An Introduction to Database Systems*. 8th ed. Boston, MA: Addison-Wesley (2004). ISBN 0-321-19784-4.
- . *An Introduction to Database Systems, Volume II*. Reading, MA:

- Addison-Wesley (1983, 1985). ISBN 0-201-14474-3.
- . *The Database Relational Model*. Boston, MA: Addison-Wesley (2001). ISBN 0-201-61294-1.
- Date, C. J., and Hugh Darwen. *Foundation for Object/Relational Databases: The Third Manifesto*. Reading, MA: Addison-Wesley (1998). ISBN 0-201-30978-5.
- Date, C. J., et al. *Temporal Data and the Relational Model*. San Francisco, CA: Morgan Kaufmann (2003). ISBN 1-55860-855-9.
- Delobel, Claude, et al. *Databases: From Relational to Object-Oriented Systems*. London, England: International Thomson Computer Press (1995). ISBN 1-850-32124-8.
- Dittrich, Klaus R., and Andreas Geppert, eds. *Component Database Systems*. San Francisco, CA: Morgan Kaufmann (2001). ISBN 1-55860-642-4.
- Dunham, Jeff. *Database Performance Tuning Handbook*. New York, NY: McGraw-Hill (1998). ISBN 0-07-018244-2.
- Egan, David, et al. *DBA's Guide to Databases on Linux*. Rockland, MA: Syngress Media (2000). ISBN 1-928994-04-0.
- Feiler, Jesse. *Database-Driven Web Sites*. San Francisco, CA: Morgan Kaufmann (1999). ISBN 0-12-251336-3.
- Fortier, Paul J. *Database Systems Handbook*. New York, NY: McGraw-Hill (1997). ISBN 0-07-021626-6.
- Gillenson, Mark L. *Database Step-by-Step*. 2nd ed. New York, NY: John Wiley & Sons (1990). ISBN 0-471-61759-8.
- Gilula, Mikhail M. *The Set Model for Database and Information Systems*. Reading, MA: ACM Press/Addison-Wesley (1994). ISBN 0-201-59379-3.
- Goldstein, Robert C. *Database Technology and Management*. New York, NY: John Wiley & Sons (1985). ISBN 0-471-88737-4.
- Goodson, John, and Robert A. Steward. *The Data Access Handbook*. Upper Saddle River, NJ: Prentice Hall (2009). ISBN 0-13-714393-1.
- Grant, John. *Logical Introduction to Databases*. Orlando, FL: Harcourt Brace Jovanovich (1987). ISBN 0-15-551175-0.
- Hackathorn, Richard. *Enterprise Database Connectivity*. New York, NY: John Wiley & Sons (1993). ISBN 0-471-57802-9.

- House, William C., ed. *Database Management*. New York, NY: Petrocelli Books (1974). ISBN 0-88405-051-3.
- Jackson, Glenn A. *Relational Database Design with Microcomputer Applications*. Englewood Cliffs, NJ: Prentice Hall (1988). ISBN 0-13-771841-1.
- Johnson, James L. *Database: Models, Languages, Design*. New York, NY: Oxford University Press (1997). ISBN 0-19-510783-7.
- Kim, Won. *Modern Database Systems*. Reading, MA: ACM Press/Addison-Wesley (1995). ISBN 0-201-59098-0.
- Korth, Henry F., and Abraham Silberschatz. *Database System Concepts*. New York, NY: McGraw-Hill (1986). ISBN 0-07-044752-7.
- Larson, James A. *Database Directions*. Upper Saddle River, NJ: Prentice Hall (1995). ISBN 0-13-290867-0.
- Lewis, Philip M., Arthur Bernstein, and Michael Kifer. *Databases and Transaction Processing*. Boston, MA: Addison-Wesley (2002). ISBN 0-201-70872-8.
- Loucopoulos, Pericles, and Roberto Zicari, eds. *Conceptual Modeling, Databases, and CASE*. New York, NY: John Wiley & Sons (1992). ISBN 0-471-55462-6.
- Martin, James. *Computer Database Organization*. Englewood Cliffs, NJ: Prentice Hall (1975). ISBN 0-13-165506-X.
- . *Managing the Database Environment*. Englewood Cliffs, NJ: Prentice Hall (1983). ISBN 0-13-550582-8.
- Mattison, Robert M. *Understanding Database Management Systems*. New York, NY: McGraw-Hill (1993). ISBN 0-07-040973-0.
- McFadden, Fred R., Jeffrey A. Hoffer, and Mary B. Prescott. *Modern Database Management*. 5th ed. Reading, MA: Addison-Wesley (1999). ISBN 0-8053-6054-9.
- Mittra, Sitansu S. *Database Performance Tuning and Optimization*. New York, NY: Springer-Verlag (2003). ISBN 0-387-95393-0.
- O'Neil, Patrick. *Database Principles, Programming, Performance*. San Francisco, CA: Morgan Kaufmann (1994). ISBN 1-55860-219-4.
- Papadimitriou, Christos. *The Theory of Database Concurrency Control*.

- Rockville, MD: Computer Science Press (1986). ISBN 0-88175-027-1.
- Parsaye, Kamran, et al. *Intelligent Databases*. New York, NY: John Wiley & Sons (1989). ISBN 0-471-50345-2.
- Pascal, Fabian. *Practical Issues in Database Management*. Reading, MA: Addison-Wesley (2000). ISBN 0-201-48555-9.
- . *Understanding Relational Databases*. New York, NY: John Wiley & Sons (1993). ISBN 0-471-58538-6.
- Piattini, Mario, and Oscar Diaz, eds. *Advanced Database Technology and Design*. Boston, MA: Artech House (2000). ISBN 0-89006-395-8.
- Podcameni, Silvio, Manfred Mittelmeir, and Michele Chilanti. *Distributed Relational Database: Cross-Platform Connectivity and Applications*. 3rd ed. Upper Saddle River, NJ: Prentice Hall (1996). ISBN 0-13-570797-8.
- Pratt, Philip J., and Joseph J. Adamski. *The Concepts of Database Management*. 2nd ed. Cambridge, MA: International Thomson Publishing (1997). ISBN 0-7600-4925-4.
- . *Database Systems Management and Design*. 2nd ed. Boston, MA: Boyd & Fraser (1991). ISBN 0-87835-579-0.
- Purba, Sanjiv, ed. *Data Management Handbook*. 3rd ed. Boca Raton, FL: Auerbach (2000). ISBN 0-8493-9832-0.
- . *Handbook of Data Management 1999*. Boca Raton, FL: Auerbach (1999). ISBN 0-8493-9976-9.
- . *High-Performance Web Databases*. Boca Raton, FL: Auerbach (2000). ISBN 0-8493-0882-8.
- Riccardi, Greg. *Principles of Database Systems with Internet and Java Applications*. Boston, MA: Addison-Wesley (2001). ISBN 0-201-61247-X.
- Rob, Peter, and Carlos Coronel. *Database Systems: Design, Implementation, and Management*. 10th ed. Cambridge, MA: Thomson Learning (2012). ISBN 978-1-111-96960-8.
- Rothstein, Michael F., and Burt Rosner. *The Professional's Guide to Database Systems Project Management*. New York, NY: John Wiley & Sons (1990). ISBN 0-471-62130-7.
- Ryan, Nick, and Dan Smith. *Database Systems Engineering*. London,

- England: International Thomson Computer Press (1995). ISBN 1-85032-115-9.
- Salemi, Joe. *Guide to Client/Server Databases*. Emeryville, CA: ZD Press (1993). ISBN 1-56276-070-X.
- Saracco, Cynthia Maro. *Universal Database Management: A Guide to Object/Relational Technology*. San Francisco, CA: Morgan Kaufmann (1998). ISBN 1-55860-519-3.
- Schur, Stephen G. *The Database Factory*. New York, NY: John Wiley & Sons (1994). ISBN 0-471-55844-3.
- Shasha, Dennis A. *Database Tuning: A Principled Approach*. Englewood Cliffs, NJ: Prentice Hall (1992). ISBN 0-13-205246-6.
- Shepherd, John C. *Database Management: Theory and Applications*. Homewood, IL: Irwin (1990). ISBN 0-256-07829-7.
- Simon, Alan R. *Strategic Database Technology*. San Francisco, CA: Morgan Kaufmann (1995). ISBN 1-55860-264-X.
- Stonebraker, Michael, ed. *Readings in Database Systems*. San Mateo, CA: Morgan Kaufmann (1988). ISBN 0-934613-65-6.
- Stonebraker, Michael, and Paul Brown. *Object-Relational DBMSs: Tracking the Next Great Wave*. 2nd ed. San Francisco, CA: Morgan Kaufmann (1999). ISBN 1-55860-452-9.
- Sweet, Frank. *Consultant's Handbook of Database Design*. Jacksonville, FL: Boxes & Arrows (1988). ISBN 0-939479-03-6.
- Thuraisingham, Bhavani M. *Data Management Systems: Evolution & Interoperation*. Boca Raton, FL: CRC Press (1997). ISBN 0-8493-9493-7.
- Ullman, Jeffrey D., and Jennifer Widom. *A First Course in Database Systems*. Upper Saddle River, NJ: Prentice Hall (1997). ISBN 0-13-861337-0.
- Widom, Jennifer, and Stefano Ceri, eds. *Active Database Systems: Triggers and Rules for Advanced Database Processing*. San Francisco, CA: Morgan Kaufmann (1996). ISBN 1-55860-304-2.
- Yang, Chao-Chih. *Relational Databases*. Englewood Cliffs, NJ: Prentice Hall (1986). ISBN 0-13-771858-6.
- Zaniolo, Carlo, et al. *Advanced Database Systems*. San Francisco, CA:

Morgan Kaufmann (1997). ISBN 1-55860-443-X.

Data Administration, Data Modeling, and Database Design

Allen, Sharon, and Evan Terry. *Beginning Relational Data Modeling*. 2nd ed. Berkeley, CA: Apress (2005). ISBN 1-59059-463-0.

Ambler, Scott W. *Agile Modeling*. New York, NY: John Wiley & Sons (2002). ISBN 0-471-20282-7.

Baca, Murtha, ed. *Introduction to Metadata, Version 3*. Los Angeles, CA: Getty Publications (2008). ISBN 978-0-89236-896-9.

Batini, Carlo, et al. *Conceptual Database Design: An Entity-Relationship Approach*. Redwood City, CA: Benjamin Cummings (1992). ISBN 0-8053-0244-1.

Bruce, Thomas A. *Designing Quality Databases with IDEF1X Information Models*. New York, NY: Dorset House (1991). ISBN 0-932633-18-8.

Carlis, John, and Joseph Maguire. *Mastering Data Modeling: A User-Driven Approach*. Boston, MA: Addison-Wesley (2001). ISBN 0-201-70045-X.

Chmura, Alan, and J. Mark Heumann. *Logical Data Modeling: What It Is and How to Do It*. New York, NY: Springer Science + Business Media (2005). ISBN 0-387-22950-7.

Chrucher, Clare. *Beginning Database Design: From Novice to Professional*. New York, NY: Apress (2007). ISBN 978-1-59059-769-9.

Codd, E. F. "Further Normalization of the Database Relational Model." In *Data Base Systems*. Courant Computer Science Symposia Series, Vol. 6. Englewood Cliffs, NJ: Prentice Hall (1972).

Durell, William R. *Data Administration: A Practical Guide to Successful Data Management*. New York, NY: McGraw-Hill (1985). ISBN 0-07-018391-0.

———. *The Complete Guide to Data Modeling*. Princeton, NJ: Data Administration, Inc. (1993). No ISBN.

English, Larry. *Improving Data Warehouse and Business Information Quality*. New York, NY: John Wiley & Sons (1999). ISBN 0-471-25383-9.

- Fleming, Candace, and Barbara von Halle. *Handbook of Relational Database Design*. Reading, MA: Addison-Wesley (1989). ISBN 0-201-11434-8.
- Halpin, Terry, and Tony Morgan. *Information Modeling and Relational Databases*. 2nd ed. Burlington, MA: Morgan Kaufmann (2008). ISBN 978-0-12-373568-3.
- Harrington, Jan L. *Relational Database Design: Clearly Explained*. 2nd ed. San Francisco, CA: Morgan Kaufmann (2002). ISBN 1-55860-820-6.
- Hay, David C. *Data Model Patterns*. New York, NY: Dorset House (1996). ISBN 0-932633-29-3.
- Hernandez, Michael J. *Database Design for Mere Mortals*. 2nd ed. Boston, MA: Addison-Wesley (2003). ISBN 0-201-75284-0.
- Hoberman, Steve. *Data Modeling Made Simple*. Bradley Beach, NJ: Technics Publications (2005). ISBN 0-9771400-0-8.
- Hogan, Rex. *A Practical Guide to Database Design*. Englewood Cliffs, NJ: Prentice Hall (1990). ISBN 0-13-690967-1.
- Inmon, W. H. *Data Architecture: The Information Paradigm*. Wellesley, MA: QED Information Sciences (1989). ISBN 0-89435-268-7.
- Kliwer, Bradley D. *Database Modeling in the PC Environment*. New York, NY: Bantam Books (1992). ISBN 0-553-08952-8.
- Lahdenmaki, Tapio, and Michael Leach. *Relational Database Index Design and the Optimizers*. Hoboken, NJ: John Wiley & Sons (2005). ISBN 0-471-71999-4.
- Lefkovits, Henry C. *IBM's Repository Manager/MVS*. Wellesley, MA: QED Information Sciences (1991). ISBN 0-89435-349-7.
- Lightstone, Sam, et al. *Physical Database Design*. San Francisco, CA: Morgan Kaufmann (2007). ISBN 978-0-12-369389-1.
- Marco, David. *Building and Managing the Meta Data Repository*. New York, NY: John Wiley & Sons (2000). ISBN 0-471-35523-2.
- Marco, David, and Michael Jennings. *Universal Meta Data Models*. New York, NY: John Wiley & Sons (2004). ISBN 0-471-08177-9.
- Modell, Martin E. *Data Analysis, Data Modeling, and Classification*. New York, NY: McGraw-Hill (1992). ISBN 0-07-042634-1.

- Mosley, Mark, et al. *The DAMA Guide to the Data Management Body of Knowledge*. Bradley Beach, NJ: Technics Publications (2009). ISBN 978-0-9771400-8-4.
- Muller, Robert J. *Database Design for Smarties: Using UML for Data Modeling*. San Francisco, CA: Morgan Kaufmann (1999). ISBN 1-55860-515-0.
- Perkinson, Richard C. *Data Analysis: The Key to Database Design*. Wellesley, MA: QED Information Sciences (1984). ISBN 0-89435-105-2.
- Pilone, Dan, with Neil Pilman. *UML 2.0 in a Nutshell*. Sebastopol, CA: O'Reilly (2005). ISBN 0-596-00795-7.
- Riordan, Rebecca M. *Designing Effective Database Systems*. Boston, MA: Addison-Wesley (2003). ISBN 0-321-29093-3.
- Rishe, Naphtali. *Database Design: The Semantic Modeling Approach*. New York, NY: McGraw-Hill (1992). ISBN 0-07-052955-8.
- Ross, Ronald G. *Entity Modeling: Techniques and Application*. Boston, MA: Database Research Group (1988). ISBN 0-941049-00-0.
- Sanders, G. Lawrence. *Data Modeling*. Danvers, MA: Boyd & Fraser Publishing Company (1995). ISBN 0-87709-066-1.
- Schmidt, Bob. *Data Modeling for Information Professionals*. Upper Saddle River, NJ: Prentice Hall (1999). ISBN 0-13-080450-9.
- Silverston, Len, et al. *The Data Model Resource Book*. New York, NY: John Wiley & Sons (2002). ISBN 0-471-15264-8.
- Simsion, Graeme C., and Graham C. Witt. *Data Modeling Essentials*. 3rd ed. San Francisco, CA: Morgan Kaufmann (2005). ISBN 0-12-644551-6.
- Stephens, Ryan K., and Ronald R. Plew. *Database Design*. Indianapolis, IN: SAMS Publishing (2001). ISBN 0-672-31758-3.
- Tannenbaum, Adrienne. *Implementing a Corporate Repository: The Model Meets Reality*. New York, NY: Wiley Professional Computing (1994). ISBN 0-471-58537-8.
- . *Metadata Solutions*. Boston, MA: Addison-Wesley (2002). ISBN 0-201-71976-2.
- Teory, Toby, et al. *Database Design: Know It All*. San Francisco, CA: Morgan Kaufmann (2002). ISBN 978-0-12-374630-6.

- Weilkiens, Tim, and Bernd Oestereich. *UML 2 Certification Guide*. San Francisco, CA: Morgan Kaufmann (2007). ISBN 978-0-12-373585-0.
- Wertz, Charles K. *The Data Dictionary: Concepts and Uses*. Wellesley, MA: QED Information Sciences (1986). ISBN 0-89435-180-X.
- Wurman, Richard Saul. *Information Anxiety*. New York, NY: Doubleday (1989). ISBN 0-385-24394-4.

Database Security, Protection, and Compliance

- Afyouni, Hassan A. *Database Security and Auditing: Protecting Data Integrity and Accessibility*. Boston, MA: Thomson (2006). ISBN 0-619-21559-3.
- Anand, Sanjay. *Sarbanes-Oxley Guide for Finance and Information Technology Professionals*. 2nd ed. Hoboken, NJ: John Wiley & Sons (2006). ISBN 0-471-78533-9.
- Ben Natan, Ron. *Implementing Database Security and Auditing*. Oxford, UK: Elsevier Digital Press (2005). ISBN 978-1-55558-334-7.
- Castano, Silvana, Mariagrazia Fugini, Giancarlo Martella, and Pierangela Samarati. *Database Security*. Wokingham, England: Addison-Wesley/ACM Press (1995). ISBN 0-201-59375-0.
- Clarke, Justin. *SQL Injection Attacks and Defense*. Burlington, MA: Syngress (2009). ISBN 978-1-59749-424-3.
- Cougias, Dorian J., et al. *Say What You Do: Building a Framework of IT Controls, Policies, Standards, and Procedures*. Lecanto, FL: Schaser-Vartan (2007). ISBN 978-0-9729039-6-7.
- Denning, Dorothy E. *Cryptography and Data Security*. Reading, MA: Addison-Wesley (1983). ISBN 0-201-10150-5.
- Fowler, Kevvie. *SQL Server Forensic Analysis*. Boston, MA: Addison-Wesley (2009). ISBN 978-0-321-54436-0.
- Ingram, Aaron, and Josh Shaul. *Practical Oracle Security*. Rockland, MA: Syngress (2007). ISBN 978-1-59749-198-3.
- IT Governance Institute. *CobiT 4.0: Control Objectives, Management Guidelines, Maturity Models*. Rolling Meadows, IL: IT Governance Institute (2005). ISBN 1-933284-37-4.

- Kenan, Kevin. *Cryptography in the Database: The Last Line of Defense*. Boston, MA: Symantec Press/Addison-Wesley (2006). ISBN 0-321-32073-5.
- Litchfield, David, et al. *The Database Hacker's Handbook: Defending Database Servers*. Indianapolis, IN: John Wiley & Sons (2005). ISBN 0-7645-7801-4.
- Mack, Mary, and Matt Deniston. *A Process of Illumination: The Practical Guide to Electronic Discovery*. Portland, OR: Discovery Center of Excellence (2004). ISBN 0-9725542-1-1.
- Olson, Jack E. *Database Archiving: How to Keep Lots of Data for a Very Long Time*. Burlington, MA: Morgan Kaufmann (2009). ISBN 978-0-12-374720-4.
- . *Data Quality: The Accuracy Dimension*. San Francisco, CA: Morgan Kaufmann (2003). ISBN 1-55860-891-5.
- Oram, Andy, and John Viega. *Beautiful Security*. Sebastopol, CA: O'Reilly (2009). ISBN 978-0-596-52748-8.
- Ottman, John B., Jr. *Save the Database, Save the World!* New York, NY: Sumo Press (2010). ISBN 978-1-4583-6368-8.
- Perry, William E. *Control in a Data Base Environment*. Wellesley, MA: QED Information Sciences (1980). ISBN 0-89435-042-0.
- Redman, Thomas C. *Data Quality: Management and Technology*. New York, NY: Bantam (1992). ISBN 0-553-09149-2.
- . *Data Quality: The Field Guide*. Boston, MA: Digital Press (2001). ISBN 1-55558-251-6.
- Reynolds, George W. *Ethics in Information Technology*. 3rd ed. Boston, MA: Course Technology (2010). ISBN 978-0-538-74622-9
- Schuler, Karen, et al. *E-discovery: Creating and Managing an Enterprisewide Program*. Burlington, MA: Syngress (2009). ISBN 978-1-59749-296-6.
- Thomas, Gwen. *Alpha Males and Data Disasters: The Case for Governance*. Orlando, FL: Brass Cannon Press (2006). ISBN 978-0-0786579-0-1.
- Thuraisingham, Bhavani. *Database and Applications Security: Integrating Information Security and Data Management*. Boca Raton, FL: Auerbach

(2005). ISBN 0-8493-2224-3.

Weill, Peter, and Jeanne W. Ross. *IT Governance: How Top Performers Manage IT Decision Rights for Superior Results*. Boston, MA: Harvard Business School Press (2004). ISBN 1-59139-253-5.

Wright, Craig, et al. *The IT Regulatory and Standards Compliance Handbook: How to Survive an Information Systems Audit and Assessments*. Burlington, MA: Syngress (2008). ISBN 978-1-59749-266-9.

Data Warehousing

Adamson, Christopher. *Star Schema: The Complete Reference*. New York, NY: McGraw-Hill (2010). ISBN 978-0-07-174432-4.

Adelman, Sid, et al. *Impossible Data Warehouse Solutions: Solutions from the Experts*. Boston, MA: Addison-Wesley (2003). ISBN 0-201-76033-9.

Barquin, Ramon, and Herb Edelstein, eds. *Planning and Designing the Data Warehouse*. Upper Saddle River, NJ: Prentice Hall (1997). ISBN 0-13-255746-0.

Devlin, Barry. *Data Warehouse from Architecture to Implementation*. Reading, MA: Addison-Wesley (1997). ISBN 0-201-96245-2.

Hackney, Douglas. *Understanding and Implementing Successful Data Marts*. Reading, MA: Addison-Wesley (1997). ISBN 0-201-18380-3.

Inmon, W. H. *Building the Data Warehouse*. 2nd ed. New York, NY: John Wiley & Sons (1996). ISBN 0-471-14161-5.

Inmon, W. H., and Richard Hackathorn. *Using the Data Warehouse*. New York, NY: John Wiley & Sons (1994). ISBN 0-471-05966-8.

Inmon, W. H., Claudia Imhoff, et al. *Building the Operational Data Store*. New York, NY: John Wiley & Sons (1996). ISBN 0-471-12822-8.

Inmon, W. H., John A. Zachman, et al. *Data Stores, Data Warehousing and the Zachman Framework*. New York, NY: McGraw-Hill (1997). ISBN 0-07-031429-2.

Kelly, Sean. *Data Warehousing in Action*. New York, NY: John Wiley & Sons (1997). ISBN 0-471-96640-1.

Kimball, Ralph. *The Data Warehouse Toolkit: The Complete Guide to*

Dimensional Modeling. 2nd ed. New York, NY: John Wiley & Sons (2002). ISBN 0-471-20024-7.

Lewis, William J. *Data Warehousing and E-Commerce*. Upper Saddle River, NJ: Prentice Hall (2001). ISBN 0-13-091154-2.

Poe, Vidette. *Building a Data Warehouse for Decision Support*. Upper Saddle River, NJ: Prentice Hall (1996). ISBN 0-13-371121-8.

Turban, Efraim. *Decision Support and Expert Systems*. Englewood Cliffs, NJ: Prentice Hall (1995). ISBN 0-02-421701-8.

Witten, Ian H., and Eibe Frank. *Data Mining: Practical Machine Learning Tools and Techniques*. San Francisco, CA: Morgan Kaufmann (2005). ISBN 0-12-088407-0.

SQL

Beighley, Lynn. *Head First SQL*. Sebastopol, CA: O'Reilly (2007). ISBN 978-0-596-52684-9.

Bhamidipati, Kishore. *SQL Programmer's Reference*. Berkeley, CA: Osborne/McGraw-Hill (1998). ISBN 0-07-882460-5.

Celko, Joe. *Data, Measurements and Standards in SQL*. San Francisco, CA: Morgan Kaufmann (2010). ISBN 978-0-12-374722-8.

———. *SQL for Smarties: Advanced SQL Programming*. 4th ed. San Francisco, CA: Morgan Kaufmann (2011). ISBN 978-0-12-382022-8.

———. *SQL Programming Style*. San Francisco, CA: Morgan Kaufmann (2005). ISBN 0-12-088797-5.

———. *Thinking in Sets: Auxiliary, Temporal, and Virtual Tables in SQL*. San Francisco, CA: Morgan Kaufmann (2008). ISBN 978-0-12-374137-0.

Cumming, Andrew, and Gordon Russell. *SQL Hacks: Tips & Tools for Digging into Your Data*. Sebastopol, CA: O'Reilly (2007). ISBN 978-0-596-52799-0.

Date, C. J., with Hugh Darwen. *A Guide to the SQL Standard*. 4th ed. Reading, MA: Addison-Wesley (1997). ISBN 0-201-96426-0.

Donahoo, Michael J., and Gregory D. Speegle. *SQL: Practical Guide for Developers*. San Francisco, CA: Morgan Kaufmann (2005). ISBN 978-0-12-220531-6.

- Faroult, Stephane. *The Art of SQL*. Sebastopol, CA: O'Reilly (2006). ISBN 0-596-00894-5.
- Groff, James R., and Paul N. Weinberg. *LAN Times Guide to SQL*. Berkeley, CA: Osborne/McGraw-Hill (1994). ISBN 0-07-882026-X.
- Gulutzan, Peter, and Trudy Pelzer. *SQL-99 Complete, Really*. Lawrence, KS: R&D Books (1999). ISBN 0-87930-568-1.
- . *SQL Performance Tuning*. Boston, MA: Addison-Wesley (2003). ISBN 0-201-79169-2.
- Harrington, Jan L. *SQL Clearly Explained*. 3rd ed. Burlington, MA: Morgan Kaufmann (2010). ISBN 978-0-12-375607-8.
- Houlette, Forrest. *Troubleshooting SQL*. Berkeley, CA: Osborne/McGraw-Hill (2001). ISBN 0-07-213489-5.
- Kline, Kevin. *SQL in a Nutshell*. 3rd ed. Sebastopol, CA: O'Reilly (2009). ISBN 978-0-596-51884-4.
- Limeback, Rudy. *Simply SQL: The Fun and Easy Way to Learn Best Practice SQL*. Collingwood, VIC, Australia: Sitepoint (2008). ISBN 978-0-9804552-5-0.
- Melton, Jim. *Advanced SQL 1999: Understanding Object-Relational and Other Advanced Features*. San Francisco, CA: Morgan Kaufmann (2003). ISBN 1-55860-677-7.
- . *Understanding SQL's Stored Procedures: A Complete Guide to SQL/PSM*. San Francisco, CA: Morgan Kaufmann (1998). ISBN 1-55860-461-8.
- Molinaro, Anthony. *SQL Cookbook*. Sebastopol, CA: O'Reilly (2006). ISBN 978-0-596-00976-2.
- Tow, Dan. *SQL Tuning*. Sebastopol, CA: O'Reilly (2004). ISBN 0-596-00573-3.

Object Orientation and Database Management

- Barry, Douglas K. *The Object Database Handbook*. New York, NY: John Wiley & Sons (1996). ISBN 0-471-14718-4.
- Brathwaite, Kenmore S. *Object-Oriented Database Design: Concepts and Application*. San Diego, CA: Academic Press (1993). ISBN 0-12-125882-

3.

Cattell, R. G. G. *Object Data Management*. Reading, MA: Addison-Wesley (1992). ISBN 0-201-53092-9.

Hughes, John G. *Object-Oriented Databases*. New York, NY: Prentice Hall (1991). ISBN 0-13-629882-6.

Kim, Wom. *Introduction to Object-Oriented Databases*. Cambridge, MA: MIT Press (1990). ISBN 0-262-11124-1.

Kroha, Petr. *Objects and Databases*. London, England: McGraw-Hill (1993). ISBN 0-07-707790-3.

Loomis, Mary. *Object Databases: The Essentials*. Reading, MA: Addison-Wesley (1995). ISBN 0-201-56341-X.

Taylor, David A. *Object-Oriented Technology: A Manager's Guide*. Reading, MA: Addison-Wesley (1990). ISBN 0-201-56358-4.

Operating Systems

Bruzzese, J. Peter, Ronald Barrett, and Wayne Dipchan. *Windows Server 2008: How To*. Indianapolis, IN: SAMS (2010). ISBN 978-0-672-33075-9.

Hassell, Jonathan. *Learning Windows Server 2003*. 2nd ed. Sepastopol, CA: O'Reilly (2006). ISBN 0-596-10123-6.

Johnson, Robert H. *MVS Concepts and Facilities*. New York, NY: McGraw-Hill (1989). ISBN 0-07-032673-8.

Nemeth, Evi, et al. *UNIX and Linux System Administration Handbook*. 4th ed. Upper Saddle River, NJ: Prentice Hall (2010). ISBN 978-0-13-148005-6.

Reiss, Levi, and Joseph Radin. *Unix System Administration Guide*. Berkeley, CA: Osborne/McGraw-Hill (1993). ISBN 0-07-881951-2.

Samson, Steven L. *MVS Performance Management*. New York, NY: McGraw-Hill (1997). ISBN 0-07-057700-5.

Young, John. *Exploring IBM's New Age Mainframes*. 5th ed. Gulf Breeze, FL: Maximum Press (1998). ISBN 978-1885068156.

Related Topics

- Abiteboul, Serge, et al. *Data on the Web: From Relations to Semistructured Data and XML*. San Francisco, CA: Morgan Kaufmann (2000). ISBN 1-55860-622-X.
- Allspaw, John. *The Art of Capacity Planning*. Sebastopol, CA: O'Reilly (2008). ISBN 978-0-596-51857-8.
- Anagol-Subbarao, Anjali. *J2EE Web Services on BEA WebLogic*. Upper Saddle River, NJ: Prentice Hall (2005). ISBN 0-13-143072-6.
- Applequist, Daniel K. *XML and SQL: Developing Web Applications*. Boston, MA: Addison-Wesley (2002). ISBN 0-201-65796-1.
- Barcia, Roland, et al. *Persistence in the Enterprise: A Guide to Persistence Technologies*. Upper Saddle River, NJ: IBM Press (2008). ISBN 978-0-13-158756-4.
- Bell, Judy Kay. *Disaster Survival Planning: A Practical Guide for Businesses*. Port Hueneme, CA: Disaster Survival Planning, Inc. (1991). ISBN 0-9630580-2-9.
- Bernstein, Philip A., and Eric Newcomer. *Principles of Transaction Processing*. San Francisco, CA: Morgan Kaufmann (1997). ISBN 1-55860-415-4.
- Ceri, Stefano, et al. *Designing Data-Intensive Web Applications*. San Francisco, CA: Morgan Kaufmann (2003). ISBN 1-55860-190-2.
- Chantico Publishing Company, Inc. *Disaster Recovery Handbook*. Blue Ridge Summit, PA: Tab Professional & Reference Books (1991). ISBN 0-8306-7663-5.
- de Guise, Preston. *Enterprise Systems Backup and Recovery: A Corporate Insurance Policy*. Boca Raton, FL: Auerbach/CRC Press (2009). ISBN 978-1-4200-7639-4.
- DeMarco, Tom, and Timothy Lister. *Peopleware: Productive Projects and Teams*. New York, NY: Dorset House (1987). ISBN 0-932633-05-6.
- Dix, Paul. *Service-Oriented Design with Ruby and Rails*. Boston, MA: Addison-Wesley (2010). ISBN 978-0-321-65936-1.
- Eddy, Sandra E. *XML in Plain English*. Foster City, CA: M&T Books (1998). ISBN 0-7645-7006-4.
- EMC Education Services. *Information Storage and Management: Storing,*

- Managing, and Protecting Digital Information*. Indianapolis, IN: Wiley Publishing (2009). ISBN 978-0-470-29421-5.
- Freedman, Daniel P., and Gerald M. Weinberg. *Handbook of Walkthroughs, Inspections, and Technical Reviews*. New York, NY: Dorset House (1990). ISBN 0-932633-19-6.
- Fronckowiak, John W. *Teach Yourself OLE DB and ADO in 21 Days*. Indianapolis, IN: SAMS Publishing (1997). ISBN 0-672-31083-X.
- Gagliardi, Gary. [*Client/Server Computing*](#). Englewood Cliffs, NJ: Prentice Hall (1994). ISBN 0-13-290784-4.
- Gane, Chris, and Trish Sarson. *Structured Systems Analysis: Tools & Techniques*. St. Louis, MO: IST Data Books (1980). ISBN 0-930196-00-7.
- Geiger, Kyle. *Inside ODBC*. Redmond, WA: Microsoft Press (1995). ISBN 1-55615-815-7.
- Ginac, Frank P. *Creating High Performance Software Development Teams*. Upper Saddle River, NJ: Prentice Hall (2000). ISBN 0-13-085083-7.
- Gunther, Neil J. *The Practical Performance Analyst*. Lincoln, NE: Authors Choice Press (2000). ISBN 0-595-12674-X.
- Gray, Jim, and Andreas Reuter. *Transaction Processing: Concepts and Techniques*. San Francisco, CA: Morgan Kaufmann (1993). ISBN 1-55860-190-2.
- Jennings, Roger. *Database Developer's Guide with Visual Basic 6*. Indianapolis, IN: SAMS Publishing (1999). ISBN 0-672-31063-5.
- Jepson, Brian. *Java Database Programming*. New York, NY: John Wiley & Sons (1997). ISBN 0-471-16518-2.
- Johnston, Tom, and Randall Weiss. *Managing Time in Relational Databases*. Burlington, MA: Morgan Kaufmann (2010). ISBN 978-0-12-375041-9.
- Kaute, Pierre Henri, Tobin Harris, Christian Bauer, and Gavin King. *NHibernate in Action*. Greenwich, CT: Manning Publications (2009). ISBN 978-1-932394-92-4.
- Khoshafian, Setrag, et al. *A Guide to Developing Client/Server SQL Applications*. San Mateo, CA: Morgan Kaufmann (1992). ISBN 1-55860-

147-3.

Koch, Richard. *The 80/20 Principle*. New York, NY: Currency/Doubleday (1998). ISBN 0-385-49170-0.

Limoncelli, Thomas A., et al. *The Practice of System and Network Administration*. Boston, MA: Addison-Wesley (2007). ISBN 978-0-321-49266-1.

Loosley, Chris, and Frank Douglas. *High Performance Client/Server*. New York, NY: John Wiley & Sons (1998). ISBN 0-471-16269-8.

Marguerie, Fabrice, Steve Eichert, and Jim Wooley. *LINQ in Action*. Greenwich, CT: Manning Publications (2008). ISBN 978-1-933988-16-0.

McClain, Gary. *OLTP Handbook*. New York, NY: McGraw-Hill (1993). ISBN 0-07-044985-6.

O'Donnell, Glenn, and Carlos Casanova. *The CMDB Imperative: How to Realize the Dream and Avoid the Nightmare*. Upper Saddle River, NJ: Prentice Hall (2009). ISBN 978-0-13-700837-7.

O'Reilly Media. *Big Data Now*. Sebastopol, CA: O'Reilly (2001). ISBN 978-1-449-31518-4.

Orfali, Robert, Dan Harkey, and Jeri Edwards. *Essential Client/Server Survival Guide*. New York, NY: Van Nostrand Reinhold (1994). ISBN 0-442-01941-6.

Patrick, Tim. *ADO.NET 4: Step by Step*. Sebastopol, CA: Microsoft Press (2010). ISBN 978-0-7356-3888-4.

Photopoulos, Constantine. *Managing Catastrophic Loss of Sensitive Data*. Burlington, MA: Syngress (2008). ISBN 978-1-59749-239-3.

Piedad, Floyd, and Michael Hawkins. *High Availability: Design, Techniques and Processes*. Upper Saddle River, NJ: Prentice Hall (2001). ISBN 0-13-096288-0.

Poelker, Christopher, and Alex Nikitin. *Storage Area Networks for Dummies*. Indianapolis, IN: Wiley Publishing (2003). ISBN 0-7645-2480-1.

Preston, W. Curtis. *Backup & Recovery*. Sebastopol, CA: O'Reilly (2007). ISBN 978-0-596-10246-3.

Pugh, Eric, and Joseph D. Gradecki. *Professional Hibernate*. Indianapolis,

- IN: Wrox (2004). ISBN 0-7645-7677-1.
- Reese, George. *Java Database Best Practices*. Sebastopol, CA: O'Reilly (2003). ISBN 0-596-00522-9.
- Rinehart, Martin. *Java Database Development*. Berkeley, CA: McGraw-Hill (1998). ISBN 0-07-882356-0.
- Scalzo, Bert, and Dan Hotka. *TOAD Handbook*. 2nd ed. Indianapolis, IN: SAMS/Developer's Library (2009). ISBN 978-0-321-64910-2.
- Sceppa, David. *Programming Microsoft ADO.NET*. Sebastopol, CA: Microsoft Press (2012). ISBN 978-0-7356-4801-2.
- Sheldon, Tom, ed. *LAN Times Guide to Interoperability*. Berkeley, CA: McGraw-Hill (1994). ISBN 0-07-882043-X.
- Strauss, Melvin J. *Computer Capacity: A Production Control Approach*. New York, NY: Van Nostrand Reinhold (1981). ISBN 0-442-26243-4.
- Sturm, Rick, Wayne Morris, and Mary Jander. *Foundations of Service Level Management*. Indianapolis, IN: SAMS Publishing (2000). ISBN 0-672-31743-5.
- SunGard Recovery Services, Inc. *Action Plan for Disaster* (1995).
- Taylor, Ed. *Demystifying SNA*. Plano, TX: Wordware (1993). ISBN 1-55622-404-4.
- . *Demystifying TCP/IP*. Plano, TX: Wordware (1993). ISBN 1-55622-400-1.
- Terplan, Kornel, and Jill Huntington-Lee. *Applications for Distributed Systems and Network Management*. New York, NY: Van Nostrand Reinhold (1995). ISBN 0-442-01873-8.
- Thornburgh, Ralph H., and Barry J. Shoenborn. *Storage Area Networks*. Upper Saddle River, NJ: Prentice Hall (2000). ISBN 0-13-027959-5.
- Toigo, Jon William. *Disaster Recovery Planning*. 2nd ed. Upper Saddle River, NJ: Prentice Hall (2000). ISBN 0-13-084506-X.
- . *The Holy Grail of Data Storage Management*. Upper Saddle River, NJ: Prentice Hall (2000). ISBN 0-13-013055-9.
- . *The Holy Grail of Network Storage Management*. Upper Saddle River, NJ: Prentice Hall (2004). ISBN 0-13-148968-2.

Walmsley, Priscilla. [*XQuery*](#). Sebastopol, CA: O'Reilly (2007). ISBN 978-0-596-00634-1.

DB2

Bond, Rebecca, et al. *Understanding DB2 9 Security*. Indianapolis, IN: IBM Press (2007). ISBN 0-13-134590-7.

Bruni, Paulo, et al. *Disaster Recovery with DB2 UDB for z/OS*. Armonk, NY: IBM Redbook (2004). ISBN 0-7384-9092-X.

Chen, Whei-Jen, et al. *High Availability and Disaster Recovery Options for DB2 on Linux, UNIX, and Windows*. Armonk, NY: IBM Redbook (2009). ISBN 0-7384-3138-9.

Chong, Raul, et al. *Understanding DB2: Learning Visually with Examples*. Indianapolis, IN: IBM Press (2005). ISBN 0-13-185916-1.

Eaton, Chris. *High Availability Guide for DB2*. Upper Saddle River, NJ: Prentice Hall (2004). ISBN 978-0-7686-8220-5.

Garvin, Curtis, and Steve Eckols. *DB2 for the COBOL Programmer Part 1*. 2nd ed. Fresno, CA: Mike Murach & Associates (1999). ISBN 1-890774-16-2.

Garvin, Curtis, and Anne Prince. *DB2 for the COBOL Programmer Part 2*. 2nd ed. Fresno, CA: Mike Murach & Associates (1999). ISBN 1-890774-03-0.

Gu, Lijun, et al. *DB2 UDB Backup and Recovery with ESS Copy Services*. San Jose, CA: IBM (2002). ISBN 0-7384-2514-1.

Mullins, Craig S. *DB2 Developer's Guide*. 6th ed. Boston, MA: IBM Press (2012). ISBN 978-0-13-283642-5.

IMS

Geller, Joseph R. *IMS Administration, Programming, and Database Design*. New York, NY: John Wiley & Sons (1989). ISBN 0-471-62185-4.

Hogan, Rex. *Diagnostic Techniques for IMS Databases*. Wellesley, MA: QED Information Sciences (1986). ISBN 0-89435-174-5.

Hubbard, George U. *IMS (DL/1) Database Organization and Performance*. New York, NY: Van Nostrand Reinhold (1986). ISBN 0-442-23583-6.

Lyon, Lockwood. *The IMS/VS Expert's Guide*. New York, NY: Van Nostrand Reinhold (1990). ISBN 0-442-23977-7.

Meltz, Dean, et al. *An Introduction to IMS*. Upper Saddle River, NJ: IBM Press/Pearson (2005). ISBN 0-13-185671-5.

MySQL

Beighley, Lynn, and Michael Morrison. *Head First PHP & MySQL*. Sebastopol, CA: O'Reilly (2009). ISBN 978-0-596-00630-3.

Bell, Charles, et al. *MySQL High Availability*. Sebastopol, CA: O'Reilly (2010). ISBN 978-0-596-80730-6.

Cabral, Sheeri K., and Keith Murphy. *MySQL Administrator's Bible*. Indianapolis, IN: John Wiley & Sons (2009). ISBN 978-0-470-41691-4.

MySQL AB. *MySQL Administrator's Guide and Language Reference*. 2nd ed. Indianapolis, IN: MySQL Press (2006). ISBN 0-672-32870-4.

Reese, George, et al. *Managing & Using MySQL*. 2nd ed. Sebastopol, CA: O'Reilly (2002). ISBN 0-596-00211-4.

Schwartz, Baron, et al. *High Performance MySQL*. 2nd ed. Sebastopol, CA: O'Reilly (2008). ISBN 978-0-596-10171-8.

Smirnova, Sveta. *MySQL Troubleshooting*. Sebastopol, CA: O'Reilly (2012). ISBN 978-1-449-31200-8.

Williams, Hugh E., and David Lane. *Web Database Applications with PHP and MySQL*. 2nd ed. Sebastopol, CA: O'Reilly (2004). ISBN 0-596-00543-1.

Yank, Kevin. *Build Your Own Database Driven Web Site Using PHP and MySQL*. 3rd ed. Collingwood, VIC, Australia: SitePoint (2004). ISBN 0-9752402-1-8.

Oracle

Bach, Martin, and Steve Shaw. *Pro Oracle Database 11g RAX on Linux*. New York, NY: Apress (2010). ISBN 978-1-4302-2959-9.

Bales, Donald. *Java Programming with Oracle JDBC*. Sebastopol, CA: O'Reilly (2002). ISBN 0-596-00088-X.

Caffrey, Melanie, and Douglas Scherer. *Oracle DBA Interactive Workbook*.

- Upper Saddle River, NJ: Prentice Hall (2001). ISBN 0-13-015742-2.
- Carpenter, Larry. *Oracle Data Guard 11g Handbook*. Berkeley, CA: McGraw-Hill/Oracle Press (2009). ISBN 978-0-07-162111-3.
- Devraj, Venkat S. *Oracle 24x7: Real-World Approaches to Ensuring Database Availability*. Berkeley, CA: Osborne/McGraw-Hill (2000). ISBN 0-07-211999-3.
- Foot, Christopher T. *OCP Instructors Guide for Oracle DBA Certification*. Kittrell, NC: Rampant Techpress (2003). ISBN 0-9744355-3-8.
- Freeman, Robert, and Matthew Hart. *Oracle RMAN 11g Backup and Recovery*. Berkeley, CA: Oracle Press/McGraw-Hill (2010). ISBN 978-0-07-162860-0.
- Harrison, Guy. *Oracle Performance Survival Guide: A Systematic Approach to Database Optimization*. Upper Saddle River, NJ: Prentice Hall (2009). ISBN 978-0-13-701195-7.
- Jesse, Scott, et al. *Oracle Database 11g Release 2 High Availability: Maximize Your Availability with Grid Infrastructure, RAC and Data Guard*. Berkeley, CA: McGraw-Hill (2011). ISBN 978-0-07-175208-4.
- Kreines, David C., and Brian Laskey. *Oracle Database Administration: The Essential Reference*. Sebastapol, CA: O'Reilly (1999). ISBN 1-56592-516-5.
- Kyte, Thomas. *Effective Oracle by Design*. New York, NY: McGraw-Hill (2003). ISBN 0-07-223065-7.
- . *Expert Oracle Database Architecture: Oracle Database 9i, 10g, and 11g Programming Techniques and Solutions*. New York, NY: Apress (2010). ISBN 978-1-4302-2946-9.
- Lawson, Christopher. *The Art and Science of Oracle Performance Tuning*. Birmingham, UK: Curlingstone (2003). ISBN 1-904347-01-0.
- Lewis, Jonathan. *Cost-Based Oracle Fundamentals*. New York, NY: Apress (2006). ISBN 1-59059-636-6.
- . *Oracle Core: Essential Internals for DBAs and Developers*. New York, NY: Apress (2011). ISBN 978-1-4302-3954-3.
- Loney, Kevin. *Oracle Database 11g: The Complete Reference*. Berkeley, CA: McGraw-Hill/Oracle Press (2008). ISBN 978-0-07-159875-0.

- Niemic, Richard. *Oracle Database 11g Release 2 Performance Tuning Tips & Techniques*. New York, NY: Oracle Press (2012). ISBN 978-0-07-178027-8.
- Owens, Kevin T. *Building Intelligent Databases with Oracle PL/SQL, Triggers, and Stored Procedures*. Upper Saddle River, NJ: Prentice Hall (1996). ISBN 0-13-443631-8.
- Price, Jason. *Java Programming with Oracle SQLJ*. Sebastopol, CA: O'Reilly (2001). ISBN 0-596-00087-1.
- Theriault, Marlene L. *Oracle8i Networking 101*. Berkeley, CA: McGraw-Hill/Oracle Press (2000). ISBN 0-07-212517-9.
- Toledo, Hugo, Jr. *Oracle Networking*. Berkeley, CA: McGraw-Hill/Oracle Press (1996). ISBN 0-07-882165-7.
- Velpuri, Rama. *Oracle Backup & Recovery Handbook*. Berkeley, CA: Osborne/McGraw-Hill (1995). ISBN 0-07-882323-4.

SQL Server

- Baird, Sean, et al. *SQL Server System Administration*. Indianapolis, IN: New Riders (1999). ISBN 1-56205-955-6.
- Bertucci, Paul. *Microsoft SQL Server High Availability*. Indianapolis, IN: SAMS (2005). ISBN 0-672-32625-6.
- Colledge, Rod. *SQL Server 2008 Administration in Action*. Greenwich, CT: Manning Publications (2010). ISBN 978-1-933988-72-6.
- Fowler, Kevvie. *SQL Server Forensic Analysis*. Boston, MA: Addison-Wesley (2009). ISBN 978-0-321-54436-0.
- Kehayias, Jonathan, and Ted Krueger. *Troubleshooting SQL Server: A Guide for the Accidental DBA*. Cambridge, UK: Red Gate (2011). ISBN 978-1-906434-77-9.
- Henderson, Ken. *The Guru's Guide to SQL Server Architecture and Internals*. Boston, MA: Addison-Wesley (2004). ISBN 0-201-70047-6.
- Horninger, Mark, et al. *The Real MCTS SQL Server 2008 Exam 70-432: Database Implementation and Maintenance Prep Kit*. Burlington, MA: Syngress (2009). ISBN 978-1-59749-420-5.
- . *The Real MCTS SQL Server 2008 Exam 70-433: Database Design*

- Prep Kit*. Burlington, MA: Syngress (2009). ISBN 978-1-59749-421-2.
- McBath, Frank. *SQL Server Backup and Recovery: Tools and Techniques*. Upper Saddle River, NJ: Prentice Hall (2001). ISBN 978-0-13-062298-3.
- Otey, Michael. *Microsoft SQL Server 2008 High Availability with Clustering & Database Mirroring*. Berkeley, CA: McGraw-Hill (2010). ISBN 978-0-07-149813-5.
- Stanek, William R. *Microsoft SQL Server 2008 Administrator's Pocket Consultant*. 2nd ed. Redmond, WA: Microsoft Press (2010). ISBN 978-0-7356-2738-3.
- Syverson, Bryan. *Murach's SQL for SQL Server*. Fresno, CA: Mike Murach & Associates (2002). ISBN 1-890774-16-2.
- Thomas, Orin, et al. *Administering Microsoft SQL Server 2012 Databases Training Kit: Exam 70-462*. Redmond, WA: Microsoft Press (2012). ISBN 978-0-7356-6607-8.

Sybase

- Garbus, Jeffrey. *Administrator's Guide to Sybase ASE 15*. Plano, TX: Jones & Bartlett (2006). ISBN 978-1-55622-360-9.
- Hitchcock, Brian. *Sybase Database Administrator's Handbook*. Upper Saddle River, NJ: Prentice Hall (1996). ISBN 0-13-357477-6.
- Kirkwood, John. *Official Sybase Internals: Designing and Troubleshooting for High Performance*. New York, NY: ITP New Media (1997). ISBN 978-1850323341.
- . *Sybase Architecture and Administration*. Hertfordshire, UK: Ellis Horwood (1993). ISBN 0-13-100330-5.
- Rozenshtein, David, et al. *Optimizing Transact-SQL: Advanced Programming Techniques*. Fremont, CA: SQL Forum Press (1995). ISBN 0-9649812-0-3.

Other Database Systems

- Anderson, J. Chris, et al. *CouchDB: The Definitive Guide*. Sebastopol, CA: O'Reilly (2010). ISBN 978-0-596-15589-6.
- Hamilton, D. D. *Inside Adabas*. Wellesley, MA: WH&O International (1995). ISBN 1-878960-03-2.

Husband, Robert E., et al. *IDMS/R Systems Desk Reference*. New York, NY: Wiley-Interscience (1987). ISBN 0-471-85236-8.

Informix Software Staff. *Evolution of the High Performance Database*. Upper Saddle River, NJ: Informix Press/Prentice Hall (1997). ISBN 0-13-594730-8.

Lumbley, Joe. *Informix Database Administrator's Survival Guide*. Upper Saddle River, NJ: Informix Press/Prentice Hall (1995). ISBN 0-13-124314-4.

White, Tom. *Hadoop: The Definitive Guide*. 3rd ed. Sebastopol, CA: O'Reilly (2012). ISBN 978-1-449-31152-0.

Worsley, John C., and Joshua D. Drake. *Practical PostgreSQL*. Sebastopol, CA: O'Reilly (2002). ISBN 1-56592-846-6.

Glossary

A

access path

The method used by the database to access data (as determined by the database optimizer).

ACID

An acronym for atomicity, consistency, isolation, and durability. Each of these four qualities is necessary for a transaction to be designed correctly.

active log

The instance of the database log that is currently in use.

Active Server Pages (ASP)

A Microsoft technology enabling programmers to develop custom code that works with Microsoft's Internet Information Server (IIS).

ad hoc query

A query constructed and executed to answer an immediate and unanticipated question or need.

ADO

Short for Microsoft ActiveX Data Objects. ADO enables client applications to access and manage data from a range of sources through an OLE DB provider.

aggregate data

Data resulting from processes that combine and summarize atomic data.

alert

The notification of an event, usually exceeding a predefined threshold.

alias

An alternative synonym for a standard name or term.

allied agent

A configured software component that works or interacts with a DBMS.

alternate key

Unique identifier for an entity instance other than the primary key.

American National Standards Institute (ANSI)

A private not-for-profit organization that coordinates the development and use of voluntary consensus standards in the United States and represents the needs of U.S. stakeholders in worldwide standardization forums (www.ansi.org).

analytics

Procedures and techniques for exploration and analysis of data to discover and identify new and meaningful information and trends.

ANSI SQL

The standard form of SQL; originally defined by ANSI nationally but now controlled by ISO.

application programming interface (API)

A published standard format for communicating with application programs.

architecture

An organized set of consensus decisions on policies, principles, services, common solutions, standards, and guidelines as well as specific vendor products and technologies used to provide information technology (IT).

archive

Database archiving is a process of removing inactive data from the operational databases on a regular basis and storing it in a different format more appropriate for long-term data retention and access when needed.

archive log

Earlier version of the database log stored for posterity.

associative entity

An entity that resolves a many-to-many relationship.

asynchronous

A style of communication in which the initiator does not wait for a reply.

atomic data

Data at the lowest available level of detail or granularity.

atomicity

The characteristic of a transaction whereby database modifications must adhere to an “all or nothing” rule. If any single part of a transaction fails, the entire transaction fails. A database management system (DBMS) must maintain atomicity despite software and hardware failures (the A in ACID).

attribute

An inherent fact, property, or characteristic describing an entity. Every attribute does one of three things: describes, identifies, or relates.

audit

A formal verification of validity, accuracy, and conformance to requirements, regulations, standards, and/or guidelines.

audit trail

Data maintained to trace activity, such as a transaction log, for purposes of recovery or audit.

authentication

The process of confirming a user’s or computer’s identity.

authorization

The granting of authority allowing a person, group, or software agent to access a resource.

automatic summary table

A database object providing automated aggregation and storage for subsequent analysis and querying. Also known as *materialized query table* or *materialized view*.

availability

The percentage of time a system or data resource is accessible compared to the time it is expected to be accessible.

B

backup

(1) The process of making a copy of data from a database to ensure its continued availability in the event of a hardware or software failure

requiring recovery of the database to restore the data. (2) The copy itself.

batch processing

A type of processing whereby computer programs are coded into jobs that are submitted for execution and run to completion without online interaction or manual intervention.

benchmark

A point of reference for measurement, comparison, and evaluation; for database systems, a benchmark typically is a point-in-time snapshot measurement for comparison with other benchmarks.

best practice

A technique, method, process, discipline, incentive, or reward generally considered to be more effective at delivering a particular outcome than by other means.

bitmap

An indexing technique using a string of zeroes and ones, or bits. For each key value of the bitmap index a separate string of zeroes and ones is stored.

Boyce-Codd normal form (BCNF)

A relation is in Boyce-Codd normal form if every determinant is a candidate key.

b-tree

A keyed, treelike index structure.

buffer pool

An area of memory set aside and used to avoid I/O operations when actual data is being read from the database. Also referred to as a *data cache*.

business intelligence (BI)

Computer-based techniques used in identifying, extracting, and analyzing business data. Common functions of BI technologies are reporting, online analytical processing (OLAP), analytics, data and process mining, complex event processing, business performance management, benchmarking, and predictive analytics.

byte

A single character of data stored electronically in 16 binary bits.

C

cache

A store of something useful; in database systems, often an area of memory for storing data, parameters, or program settings.

call-level interface (CLI)

Defines how a program should send SQL queries to the DBMS and how the returned data should be handled by the application in a consistent way.

candidate key

An attribute or set of attributes that can be used to uniquely identify an occurrence of the entity. Each entity may have one or more candidate keys. One of these candidate keys is selected as the table primary key.

cardinality

The number of occurrences that may exist between a pair of entities. Another way of looking at cardinality is as the number of entity occurrences applicable to a specific relationship. Sometimes the term *degree* is used instead of cardinality. An alternate usage of the term *cardinality* within the realm of database administration is a database statistic used by the relational optimizer defining the number of occurrences of a value within a column (or set of columns).

Cartesian product

The result of a Cartesian join, occurring when every row of one table is joined to every row of another table (or itself) without the use of join predicates.

centralized processing

The use of computer assets consolidated in a single computer (or cluster of coupled computers) and accessed using dumb terminals.

certification

The process by which an organization grants recognition to an individual, organization, process, service, or product that meets certain established criteria.

change data capture

The process of capturing changes made to a production data source;

typically used in data warehousing environments.

cipher

An algorithm for performing encryption or decryption.

CLI

See [call-level interface](#).

client/server computing

A multitier approach to computing where applications access corporate shared resources over a network connection.

cloud computing

The delivery of computing as a service. Cloud computing applications rely on a network (typically the Internet) to provide users with shared resources, software, and data.

clustering

(1) The condition whereby data is physically ordered contiguously by a specified key (usually implemented by means of an index). (2) The use of multiple, “independent” computing systems working together to form what appears to users as a single highly available system.

COBIT (Control Objectives for Information and Related Technology)

Standards for IT governance published by the IT Governance Institute; see www.aisca.org/cobit.

CODASYL (Conference on Data Systems Languages)

A consortium formed in 1959 to guide the development of a standard programming language that could be used on many computers. This effort led to the development of the network (navigational) data model for DBMS and COBOL, among other standards.

column

A data attribute as implemented in a relational database as a vertical component of a table, similar to a field in a flat file record.

COM (Common Object Model)

Microsoft’s programming specification for object interoperability through sets of predefined routines called interfaces.

commit

SQL statement that concludes a unit of work.

compliance

The process of conforming, completing, performing, or adapting actions to meet the rules, demands, or wishes of another party. Commonly used when discussing conformance to external government or industry regulations.

composite key

A key with more than one attribute.

conceptual data model

A high-level data model identifying major entities and relationships, not fully attributed and therefore not necessarily normalized.

concurrency

The condition when more than one user has access to a specific set of data at the same time. Concurrency is also used to describe a database management system's capability to handle simultaneous queries against a single set of tables.

config.ora

A file associated with the Oracle client that specifies certain defaults, files, and directory names for the Oracle client.

consistency

The degree to which data values are equivalent across redundant databases. With regard to transactions, consistency refers to the state of the data both before and after the transaction is executed. A transaction maintains the consistency of the state of the data. In other words, after a transaction is run, all data in the database is "correct" (the C in ACID).

constraint

(1) A restriction on a business action and the resulting data. (2) The database mechanism for enforcing such.

contention

A conflict over access to a shared resource such as memory, disk storage, a database object, or a data page.

CORBA (Common Object Request Broker Architecture)

The Object Management Group's vendor-independent architecture and

infrastructure for object-based programming interoperability.

cursor

A moving placement or pointer that indicates a position, used in programs to overcome the impedance mismatch between SQL (set-at-a-time processing) and application programs (record-at-a-time processing).

D

data

Facts represented as text, numbers, graphics, images, sound, or video (with no additional defining context); the raw material used to create information.

database

An organized collection of data, usually in digital form.

database administrator

An individual responsible for the installation, configuration, administration, monitoring, and maintenance of the DBMS and its physical databases.

database driver

Software that enables applications to request and retrieve data over the network from the database server using database protocol packets.

database gateway

Software required to allow clients to access data stored on database servers over a network connection.

database management system (DBMS)

A software package that can be used to control the creation, maintenance, and use of a database. It allows organizations to conveniently develop databases for various applications.

database writer (DBWR)

An Oracle process that writes data from the data cache contained in memory out to the physical disk files.

data breach

The release of secure information to an untrusted environment. Other terms

for this occurrence include *unintentional information disclosure*, *data leak*, and *data spill*.

data cache

An area of memory set aside and used to avoid I/O operations when actual data is being read from the database. Also referred to as a *buffer pool*.

data cleansing

The act of detecting and removing and/or correcting data in a database. Also called *data scrubbing*.

Data Control Language (DCL)

SQL statements used for controlling access to data and data structures, e.g., GRANT and REVOKE.

Data Definition Language (DDL)

SQL statements used for creating and modifying data structures, e.g., CREATE, ALTER, and DROP.

data dictionary

Software coupled with a data store for managing data definitions.

Data Manipulation Language (DML)

SQL statements used for accessing and modifying data, e.g., INSERT, UPDATE, DELETE, and SELECT.

data stewardship

Responsibility and accountability for the actions taken upon a defined set of data, including the definition of the consumers of the data. A data steward is not necessarily the data owner.

data type

A classification identifying one of various types of data. Each column in a physical database design must have a data type assigned. Examples include integer, character, etc.

data warehouse

A database used for reporting and analysis.

DBA

See [database administrator](#).

DBCC

Sybase utility for consistency and integrity checking.

DCL

See [Data Control Language](#).

DDL

See [Data Definition Language](#).

deadlock

A specific type of locking problem when concurrent processes are competing for locks. For example, program1 holds a lock on A and is waiting for a lock on B; program2 holds a lock on B and is waiting for a lock on A.

decryption

The reverse of encryption; making the encrypted information readable again (usually requiring a key).

delete

A SQL statement for removing rows from a relational database.

denormalization

Undoing the effect of normalization; the process of putting one fact in numerous places in the database.

design reviews

A series of meetings wherein all aspects of the database and application code are reviewed for efficiency, effectiveness, and accuracy.

dimension

In a data warehouse, a data element that categorizes each item in a data set into nonoverlapping regions.

dirty data

Data that is incorrect, out-of-date, redundant, incomplete, or formatted incorrectly.

disaster

Any event that has a small chance of transpiring, a high level of uncertainty, and a potentially devastating outcome.

disaster recovery

The process and planning involved in recovering a data center from a disaster.

distributed processing

The use of multiple computers to fulfill a service request.

DML

See [Data Manipulation Language](#).

domain

Defines the universe of valid values for a data element.

DSNZPARM

System configuration parameters for mainframe DB2 subsystems.

DSN1COPY

A DB2 utility enabling database files to be copied outside the control of the DBMS.

durability

Refers to the impact of an outage or failure on a running transaction. A durable transaction will not impact the state of data if the transaction ends abnormally. The data will survive any failures (the *D* in ACID).

dynamic SQL

SQL queries that are not preprocessed and whose access paths are determined at run time right before execution.

E**encapsulation**

(1) A language mechanism for restricting access to some of an object's components. (2) A language construct that facilitates the bundling of data with methods (or other functions) operating on that data.

encryption

The process of transforming data using an algorithm (called a *cipher*) to make it unreadable to anyone except those possessing special knowledge, usually referred to as a key.

enterprise data model

A single data model that comprehensively describes the data needs of the entire organization.

entity

Something that exists and is capable of being described. It is a person, place, thing, concept, or event about which your organization maintains facts.

entity integrity

The most basic level of data integrity provided by relational databases stating that each occurrence of an entity must be uniquely identifiable.

entity-relationship (E/R) diagram

An E/R diagram graphically depicts the entities and relationships of a data model.

ETL

Shorthand for extract/transform/load, describing a category of tools used for populating data warehouses.

exclusive lock

The type of lock taken by the DBMS when data is modified to prohibit other processes or users from reading or modifying the same data.

explain

Command (DB2, Oracle) used to externalize the access paths chosen by the relational optimizer.

expression

An instruction to execute something that will return a value.

F**fact table**

A central table in a dimensional model that contains numerical measures and key facts relating to dimension tables.

fallback

The process of returning to a prior release of software after having updated to a new release.

fat client

A client/server processing model in which significant functionality resides on the client computer.

field

An attribute implemented on a screen or flat file record.

fifth normal form (5NF)

Specifies that every join dependency for the entity must be a consequence of its candidate keys.

first normal form (1NF)

An entity is in first normal form if and only if all underlying domains contain atomic values only.

FISMA

Federal Information Security Management Act, also known as the E-Government Act. Governmental regulation stating that federal agencies, contractors, and any entity that supports them must maintain security commensurate with potential risk.

fixpack

A grouping of software patches that can be applied as a whole (as opposed to individually) to correct bugs and security vulnerabilities.

flat file

Plain text file wherein each line of the file holds one record, typically with fields separated by delimiters, such as commas or tabs.

foreign key

An attribute or set of attributes that identify relationships between entity occurrences.

fourth normal form (4NF)

An entity is in fourth normal form (4NF) if and only if it is in 3NF and has no multiple sets of multivalued dependencies.

function

A predefined algorithm supported by the DBMS for converting, manipulating, or calculating data based upon input. Functions can be used within SQL statements any place an expression can be used.

G

GDG (generation data group)

A type of file structure used on mainframe systems in which a group of data sets are related to each other (chronologically and functionally) and share a unique data set name. Every GDG data set has a generation and version number assigned to it.

gigabyte (GB)

A billion bytes of storage; 1,000 megabytes (MB).

GLB

Gramm-Leach-Bliley Act, also known as the Financial Modernization Act of 1999. A federal law enacted in the United States to control the ways that financial institutions deal with the private information of individuals.

governance

Consistent management, cohesive policies, guidance, processes, and decision rights for a given area of responsibility. For example, corporate governance can involve policies on privacy, internal investment, and the use of data.

H

hashing

A technique that transforms a key value via an algorithm to a physical storage location to enable quick direct access to data. The algorithm is typically referred to as a *randomizer*, because the goal of the hashing routine is to spread the key values evenly throughout the physical storage.

heterogeneous

Relating to the uniformity or lack thereof in a system. A system that is homogeneous is uniform in composition or character; one that is heterogeneous lacks uniformity in at least one of these qualities. For example, an application that uses both Oracle and DB2 would be said to be heterogeneous.

hierarchical

An arrangement of items (objects, names, values, categories, etc.) in which

the items are represented as being “above,” “below,” or “at the same level as” one another. In database parlance, a DBMS is said to be hierarchical if the relationships among data in the database are established such that one data item is present as the subordinate of another one.

HIPAA (Health Insurance Portability and Accountability Act)

Governmental regulation specifying that health care providers must protect individuals’ health care information.

HTML (HyperText Markup Language)

The primary markup language for Web pages.

I

IDE (Integrated Development Environment)

A software application that provides comprehensive facilities to computer programmers for software development.

identity property

A feature supported by several popular relational DBMS products. When the identity property is assigned to a column, the DBMS treats that column as a store for automatically generated incremental values. The user does not provide values for the column when rows are inserted into the table; the DBMS generates them instead.

incremental backup

A database backup containing only the data that has changed since the last full backup or incremental copy was made.

index

A physical database object used to quickly access data in a database.

information

Data with additional context in the form of metadata, including definition and relationships between data and possibly other information. Data in context with metadata makes information.

infrastructure

The equipment, systems, software, and services used in common across an organization, regardless of mission/program/project.

insert

The SQL statement used to add data to a table.

instance

(1) An occurrence of an entity. (2) The implementation of a database server (e.g., an Oracle instance).

instantiate

To create an instance of a software object or database row.

intent lock

A type of lock placed on higher-level database objects when a user or process acquires locks on data pages or rows. An intent lock stays in place for the life of the lower-level locks.

isolation

A property of transactions meaning that multiple transactions can run at the same time. Any transactions running in parallel have the illusion that there is no concurrency. In other words, it appears that the system is running only a single transaction at a time. To achieve isolation, a locking mechanism is required (the *I* in ACID).

isolation level

Specifies the locking behavior for a transaction or statement. The higher the isolation level, the stricter the locking protocol becomes. Examples include UNCOMMITTED READ, COMMITTED READ, REPEATABLE READ, and SERIALIZABLE.

J**Java Database Connectivity (JDBC)**

A call-level interface (CLI) to allow Java programs to interact with databases. Can be thought of as ODBC for Java programs.

join

An operation in which the data from two tables is combined into a larger results table based on shared data values in each table.

JVM

Java virtual machine.

J2EE (Java 2 Platform, Enterprise Edition)

A set of coordinated specifications and practices that together enable solutions for developing, deploying, and managing multitier enterprise Java applications.

K

key

The attributes that identify entity occurrences and define relationships between entities.

kilobyte (KB)

A thousand bytes of storage.

knowledge

The fact or condition of knowing something with familiarity gained through experience or association. Knowledge adds understanding and retention to information.

L

label-based access control (LBAC)

A form of fine-grained database security allowing control over who can read and write to individual rows and columns.

lock

A DBMS function used to ensure the integrity of data. When a database resource is locked by one process, another process is not permitted to change the locked data. Locking is necessary to enable the DBMS to facilitate the ACID properties of transaction processing.

lock escalation

The process of increasing the lock granularity for a process or program.

log

A collection of records describing the sequence of events that occur during DBMS execution to be used for database recovery in the event of a DBMS failure.

logical data model

A detailed data model consisting of fully normalized entities with all attributes defined. Furthermore, the domain or data type of each attribute must be defined, along with a primary key for each entity. A logical data model should be a complete document from which a physical database can be developed.

log off-loading

The process of moving active log data to an archive log (either on disk or on tape) when the active log file fills up and the log records must be retained.

log writer (LGWR)

An Oracle process that manages the redo log buffer.

M

many-to-many relationship

A link between two entities in which the cardinality of both sides of the relationship is multiple.

megabyte (MB)

A million bytes of storage; 1,000 kilobytes (KB).

memory

Data storage in the form of chips.

messaging queuing

A form of data movement whereby data, in the form of messages, is placed onto a queue by one application or process and is read from a queue by another application or process.

metadata

The simplest definition of metadata is “data about data.” To be a bit more precise, metadata describes data, providing information such as type, length, textual description, and other characteristics.

middleware

Software that connects two otherwise separate applications.

multidimensional data

Cubed data used in OLAP and data warehousing; the data consists of

numeric facts categorized by dimensions for analytical querying.

N

NAS (network-attached storage)

Refers to storage that can be accessed directly from the network. With NAS, hosts or client systems can read and write data over a network interface.

.NET

A Microsoft framework for program development providing a comprehensive platform for the construction, deployment, and management of applications.

normalization

The process of organizing data to minimize redundancy and remove ambiguity. In simple terms, normalization is the process of identifying the one best place each fact belongs.

null

Represents missing or unknown information at the attribute (or column) level. If an attribute (or column) “value” can be null, it can mean one of two things: The attribute/column is not applicable for certain occurrences of the entity, or the attribute applies to all entity occurrences, but the information may not always be known. It could be a combination of these two situations, too.

O

object

A structure consisting of data and methods together with their interactions. An object is not just data, but the encapsulation of its state (data) and behavior (methods).

object oriented

A programming paradigm that relies on objects (see [object](#)) to design applications and computer programs.

object relational

Originally intended as a merging of object-oriented technology with

relational, but in general practice simply refers to a DBMS that supports large multimedia data types and gives users the ability to define their own data types and functions.

occurrence

An instance of an entity.

ODBC (Open Database Connectivity)

A call-level interface (CLI) for interacting with databases. ODBC provides routines to allocate and deallocate resources, control connections to the database, execute SQL statements, obtain diagnostic information, control transaction termination, and obtain information about the implementation.

one-to-many relationship

A link between two entities in which the cardinality of one side of the relationship is one and the other is multiple.

one-to-one relationship

A link between two entities in which the cardinality of both sides of the relationship is one.

online analytical processing (OLAP)

A type of computer processing that provides analysis of data stored in a database. OLAP tools enable users to analyze different dimensions of multidimensional data.

online transaction processing (OLTP)

A type of computer processing in which the computer responds immediately to user requests. Each request is a transaction. The opposite of transaction processing is batch processing.

OO

Object oriented or object orientation.

Open Database Connectivity (ODBC)

A standard database access technology developed by Microsoft, the purpose of which is to allow accessing any DBMS from any application, regardless of which DBMS is managing the data.

operating system

The control program that enables a computer to function. An operating

system performs basic tasks, such as recognizing keyboard input, sending output to the screen, keeping track of files and directories on disk, and controlling peripheral devices such as disk drives and printers.

operational data store (ODS)

A database designed to integrate data from multiple sources for additional operations on the data.

optical disc

A flat disc that encodes binary data (bits) in the form of pits (binary value of 0 or off, due to lack of reflection when read) and lands (binary value of 1 or on, due to a reflection when read) on a special material (often aluminum) on one of its flat surfaces. Examples include CDs and DVDs.

optimizer

The component of a relational database system responsible for analyzing SQL queries and producing optimal access paths for retrieving data from the database.

optionality

In data modeling, whether relationships are mandatory or optional is commonly referred to as the optionality of the relationship.

ORM (object-relational mapping)

Whereby an object's attributes are stored in one or more columns of a relational table.

P

Pareto Principle

The general observation that a small amount of effort can derive a great amount of rewards. Also known as the 80/20 rule because it often is stated as 80 percent of the results come from 20 percent of the effort.

partition

(1) A subset of a database file. (2) A physical operating system storage area that bypasses the file system.

password

A secret word or string of characters used for authentication, to prove

identity or to gain access to a resource.

PCI DSS (Payment Card Industry Data Security Standard)

Industry regulation developed by the major credit card companies to help prevent credit card fraud, hacking, and other security issues.

physical data model

The logical data model transformed into a physical implementation using a specific DBMS product (e.g., DB2, Oracle, SQL Server, etc.).

predicate

A statement that can be evaluated as true or false, such as a WHERE clause of a SQL statement.

primary key

The attribute, or set of attributes, that is used to uniquely identify an occurrence of the entity. Each entity must have one and only one primary key.

privilege

An action or capability, such as table access, that can be granted to users of the DBMS.

proactive

Engaging in corrective behavior or activity in advance of experiencing a problem.

process monitor (PMON)

An Oracle background process that performs cleanup when a user process fails with an error condition.

program global area (PGA)

Basic memory structure for Oracle processes used as a work area.

propagation

The distribution of data from one or more source databases to one or more target databases, according to predefined rules.

Q

query

(1) A SQL SELECT statement written and issued against a database. (2) To request data from a database.

R

RAID (Redundant Array of Independent [or Inexpensive] Disks)

A category of disk drives that employ two or more drives in combination to deliver fault tolerance and improved performance.

reactive

Engaging in corrective behavior or activity only after experiencing a problem.

read efficiency

A formula for calculating the efficiency of a data cache (or buffer pool). It is calculated as the number of actual I/Os performed subtracted from the total number of data requests, then divided by the total number of data requests.

record

The physical representation of data about an instance in a flat file.

recovery

The restoration of a database to its state as of a previous point in time, usually in response to a hardware or software failure.

recovery time objective (RTO)

An expression of the amount of time a business will tolerate the computing system (hardware, software, DBMS, services) to be offline.

redo log files

A record of changes made to data used in recovery scenarios (Oracle).

referential constraint

A specification whereby the DBMS forces the values of a foreign key of one table to contain either null, or existing primary key values of another (or perhaps the same) table.

referential integrity (RI)

Through the specification of appropriate referential constraints, RI guarantees that an acceptable value is always in each foreign key column.

regulatory compliance

The act of meeting the requirements of governmental legislation or industry mandates.

relational closure

The feature of relational database systems whereby every operation performed on a relational database operates on a table (or set of tables) and results in another table.

relational model

A database model based on first-order predicate logic, first formulated and proposed in 1969 by Edgar F. Codd (IBM). Most of today's popular DBMS offerings are loosely based on the relational model.

relationship

Defines how different entities are associated with each other. A relationship is defined by the keys: the primary key in the parent entity and the foreign key in the dependent entity.

release

A release represents a unique state of software. A release is typically minor, with fewer changes and not as many new features as a version.

reliability

The ability for a component (server, application, database, etc.) or group of components to consistently perform its functions.

remote procedure call (RPC)

A mechanism for invoking a service on another platform.

replication

The copying of data from a data source to one or more target environments based on rules.

repository

A store of information about the data assets of an organization.

RMAN

Database system software for managing the backup and recovery of Oracle data.

role

(1) Name used to refer to the logical set of related responsibilities assignable to a person or organization, and to parties with these assigned responsibilities. (2) A database security mechanism used to grant one or more preassigned privileges to a user.

roll back

To undo the database statements performed prior to a commit of the transaction.

row

A set of column values describing one logical instance in a table.

RTO

See [recovery time objective](#).

S**SAN (storage area network)**

Generally refers to an interconnected network of storage devices.

Sarbanes-Oxley (SOX)

The Sarbanes-Oxley Act, officially known as the U.S. Public Company Accounting Reform and Investor Protection Act of 2002. Governmental regulation with the goal of reducing fraud and conflicts of interest, improving disclosure and financial reporting, and strengthening confidence in public accounting.

scalability

The ability to increase or decrease size or capability in cost-effective increments with minimal impact on the unit cost of business and the procurement of additional services.

scan

The process of reading all the rows in a table (or tablespace) to satisfy a query (as opposed to indexed access).

schema

(1) The logical or physical definition of data elements, physical characteristics, and relationships. (2) The diagram representing a logical

data model.

second normal form (2NF)

An entity is in second normal form if and only if it is in first normal form and every non-key attribute is fully dependent on the key.

security

The prevention of unauthorized access to a database and its data, and to applications that have authorized access to databases.

select

A SQL statement for specifying data retrieval operations for a relational database.

service-level agreement (SLA)

A contracted guarantee of service delivery for a program, transaction, service, or workload.

service-level management (SLM)

The disciplined, proactive methodology and procedures used to ensure that adequate levels of service are delivered to all IT users in accordance with business priorities and at acceptable cost.

service-oriented architecture (SOA)

An application architecture in which all functions, or services, are created with invocable interfaces that are called to perform business processes.

shared disk

A form of clustering in which all of the connected systems share the same disk devices.

shared lock

The type of lock taken by the DBMS when data is read with no intent to update it.

shared nothing

A form of clustering in which each system has its own private resources (memory, disks, etc.).

SHOWPLAN

Command (SQL Server, Sybase ASE) used to externalize the access paths

chosen by the relational optimizer.

Simple Object Access Protocol

See [SOAP](#).

SOAP (Simple Object Access Protocol)

A protocol specification for exchanging structured information in a Web services implementation.

solid-state drive (SSD)

A data storage device that uses solid-state memory to store persistent data.

SOX

See [Sarbanes-Oxley](#).

SQL

Originally an initialization of Structured Query Language but now stands on its own. Pronounced “sequel” or “ess-cue-el,” SQL is the de facto standard language for accessing relational databases.

SQL injection

A form of Web hacking whereby SQL statements are specified in a Web form to expose data to the attacker.

standby database

An identical copy of an online production database that is close to being up-to-date in terms of its data content. When a failure occurs, control is transferred to the standby database, which is then opened as the online production database to allow normal activity to continue.

star schema

A common form of a dimensional data model, where a fact table is directly linked by foreign keys to several dimension tables.

static SQL

SQL queries that are preprocessed and whose access paths are determined during the bind procedure, prior to execution.

stored procedure

A precompiled routine or program stored within the DBMS and executable using DBMS commands.

system global area (SGA)

A group of shared memory structures for Oracle database instances that contain data and control information.

system monitor (SMON)

An Oracle background process that provides instance recovery during start-up.

T**table**

A two-dimensional arrangement of data into rows and columns.

tablespace

A structure or database object used by some DBMS products for storing and organizing database tables.

TCO

See [total cost of ownership](#).

temporal

Having to do with time.

terabyte (TB)

A trillion bytes of storage; 1,000 gigabytes (GB).

thin client

A client/server processing model in which only limited functionality resides on the client computer.

third normal form (3NF)

An entity is in third normal form if and only if it is in second normal form and every non-key attribute is nontransitively dependent on the primary key.

tier

A level of separation of computing responsibility.

time-out

After a process waits for the predetermined amount of time (usually dictated by a system parameter), it is canceled. Time-outs typically are

associated with waiting for a lock to be granted.

total cost of ownership (TCO)

The cost to own, implement, and maintain a product throughout its life.

TPC

The Transaction Processing Performance Council, or TPC, is an independent, not-for-profit organization that manages and administers performance benchmark tests.

traffic

In network operations, the number of packets traversing the network.

transaction

An atomic unit of work with respect to recovery and consistency.

trigger

An event-driven specialized procedure that is attached to database tables; typically implemented to support data integrity requirements.

two-phase COMMIT (2PC)

A feature of transaction processing systems that enables a database to be returned to the pretransaction state if an error condition occurs.

tuple

A set of elements; the formal mathematical term for a row in a relational table.

U

union

A SQL function that concatenates the results of two SELECT statements with consistent column structures into a single answer set.

unit of work (UOW)

A series of processes, instructions, and messages that, when executed, guarantee data integrity. It is important that transactions operate in terms of units of work whereby either all, or none, of the processes are performed to completion.

update

(1) A SQL statement for modifying data within a relational database. (2) A general term referring to any change made to a database.

update lock

The type of lock taken by the DBMS when data must be read first before it is modified. The update lock indicates that the data may be modified or deleted in the future. If the data is actually modified or deleted, the DBMS will promote the update lock to an exclusive lock.

user

Person or role recognized and authorized to access a particular application or database resource.

user-defined data type (UDT)

Configurable data type added to the DBMS to be used like built-in data types.

user-defined function

Extensible code added to the DBMS to be callable like built-in database functions.

V

value

A data abstraction assigned to a single attribute representing a fact.

version

A version represents a unique state of software. A new version of software is a major concern, with many changes and new features. Contrast with **release**.

view

A virtual table, defined as a SQL SELECT statement, to provide a subset of data from one or more tables.

W

workload

The amount of work performed or capable of being performed by a system, usually within a specific period of time.

X

XML (eXtensible Markup Language)

A markup language, like HTML, that is based on SGML (Standard Generalized Markup Language). HTML uses tags to describe how data appears on a Web page. It is commonly used in Web applications to describe data.

XQuery

A query language for retrieving data from an XML document or XML database.

Z

zap

Utility or command for modifying data at the byte (or bit) level; usually used to resolve system problems such as incorrect pointer values.

z Series

IBM's mainframe computing hardware.

Index

NOTES:

Page numbers ending with an italic f (e.g. 192*f*) indicate tables or figures.

Numbers containing a lowercase n (e.g. 301n or 432n.5) indicate a footnote.

24/24 availability, [270–271](#)

80/20 (Pareto) rule, [302](#)

99.999% availability, [273–274](#)

100 Year Archive Requirements Survey, [503n](#)

A

Abbreviations, standards for, [96](#)

Absolute positioning, [383–385](#)

Access (Microsoft), [767](#)

Access paths, [187](#)

Accessibility, DBA to coworkers, [745](#)

Accessing data. *See* [Data access](#).

Accounts. *See* [Logins](#).

ACID (atomicity, consistency, isolation, and durability) properties, [205–206](#)

Acquire/release specification, [218](#)

Action Corporation, [764](#)

Active databases, [426](#)

Active metadata sources, [690](#)

Actuate, [773](#)

Adabas (Software AG), [764](#)

Adaptive Ltd., [772](#)

Adaptive Server Enterprise, 64

AD/Cycle (IBM), [695](#)

Adding objects. *See* [ALTER statements](#).

Adelphia, [485](#)

Adjectives as

 attributes, [116](#), [124](#)

- entities, [115](#)
- Adjust tables, data page layouts, [592](#)
- ADLC (application development life cycle), [9–10](#), [12](#)
- ADO.NET, [194–195](#)
- Advanced analytics, [269–270](#)
- AES (Advanced Encryption Standard), [472](#)
- Agile Modeling, [783](#)
- “Airline magazine” syndrome, [266](#)
- ALL privileges, [458–459](#)
- Allen Systems Group, [770](#), [772](#)
- Allied agents, [321–322](#)
- Allocation pages, data page layouts, [589](#)
- Allocation units, data page layouts, [589](#)
- ALLOW UPDATES parameter, [345](#)
- ALTER statements
 - changing database structures, [250–252](#), [252–253](#)
 - changing management, [701–703](#)
 - limitations, [252–253](#)
 - purpose of, [250–252](#)
- ALTER TABLE statements, [436–437](#)
- AlwaysOn features, [285](#)
- Analytical processing *versus* transaction processing, [638–640](#)
- Analytics
 - advanced, [269–270](#)
 - availability requirements, [268–270](#)
 - benefits of, [269](#)
 - DBA rule of thumb, [741–742](#)
 - tools for, [721–724](#)
- ANSI Web site, [782](#)
- APIs, SQL, [192–193](#)
- Applets, [196–197](#)
- Application DBA, [34–35](#)
- Application development life cycle (ADLC), [9–10](#), [12](#)

Application Security, [770](#)

Application servers, [209–210](#), [664](#)

Application time, [179–180](#)

Applications. *See also* [Database applications](#).

- availability problems, [279–280](#)
- backing up, [516–517](#)
- code, design review, [238](#)
- complexity, upgrading the DBMS, [88](#)
- criticality of, ranking, [562–563](#)
- DBA staffing requirements, [38–39](#)
- development standards, [100](#)
- infrastructure design, [193–194](#)
- integration, [624–625](#)
- Java, [196–197](#)
- performance, [312](#), [711–713](#). *See also* [Relational optimization](#); [SQL tuning](#).

Approach (Lotus), [767](#)

Aquafold, [770](#)

Archiving database logs, [77–78](#), [339](#), [529](#). *See also* [Backup](#).

Archiving databases

- 100 Year Archive Requirements Survey, [503n](#)
- data life cycle, [499–500](#)
- definition, [500](#)
- e-discovery, effects on DBA, [506–507](#)
- hardware independence, [503–505](#)
- hardware obsolescence, [504](#)
- overview, [500–505](#)
- versus* purging databases, [501](#)
- requirements, [503–505](#)
- software independence, [503–505](#)
- system components, [505–506](#)

Archiving databases, data retention

- overview, [498](#)

scope, determining, [501–503](#)

Associative entities, [127](#)

Atomicity

definition, [758](#)

transactions, [205–206](#)

Atomicity, consistency, isolation, and durability (ACID) properties, [205–206](#)

Attributes

adjectives as, [116](#), [124](#)

data types, [116](#)

definition, [115](#)

discovering, [124–125](#)

domains, [116](#)

missing values, [119–120](#)

naming conventions, [116–119](#)

nouns as, [124](#)

nulls, [119–120](#)

prepositional phrases as, [124](#)

purpose, [115–116](#)

transforming to columns, [142–143](#)

values, [119](#)

Auditing

security, [477–478](#)

tools for, [717–719](#)

Auditing databases

common questions, [495](#)

comprehensive methods, [494](#)

data access tracking, [490–493](#)

guidelines, [492–493](#)

log-based auditing, [493–495](#)

network sniffing, [494–495](#)

noninvasive methods, [494](#)

overview, [490–493](#)

parsing database logs, [493–495](#)

- privileged users, [495–496](#)
- regulatory requirements, [491](#)
- selective methods, [494](#)
- tapping requests, [494–495](#)
- techniques for, [493–495](#)
- trace-based auditing, [493–495](#)

Authentication, [452](#)

Authority. *See also* [Privileges](#); [Security](#).

- LBAC (label-based access control), [463–465](#)

Authority, granting privileges

- centralized administration, [457](#)
- database object privileges, [459](#)
- DCL (Data Control Language), [456–457](#)
- decentralized administration, [457](#)
- overview, [456–457](#)
- procedure privileges, [460](#)
- program privileges, [460](#)
- to PUBLIC authority, [460–461](#)
- system privileges, [459–460](#)
- table privileges, [458–459](#)
- types of privileges, [457–458](#). *See also specific types.*

Authority, revoking privileges

- cascading REVOKEs, [462](#), [468](#)
- chronology and REVOKEs, [462–463](#)
- overview, [461](#)

Authorization. *See also* [Privileges](#); [Security](#).

- availability problems, [280](#)
- database administration tasks, [24–25](#)
- database administrator, [467](#)
- database maintenance, [467](#)
- for groups, [468](#)
- operations control, [467](#)
- for roles, [466](#), [468](#)

security administrator, [467](#), [468](#)

system administrator, [467](#)

Automatic summary tables, [652–653](#)

Automation

change management, [245](#)

DBA functions, for availability, [290–291](#)

DBA rule of thumb, [737–739](#)

Autonomy, distributed databases, [626](#)

Availability. *See also* [Downtime](#).

“airline magazine” syndrome, [266](#)

change management, [246](#)

components of, [267–268](#)

database administration tasks, [24](#)

DBA staffing requirements, [38](#)

definition, [267](#)

versus downtime, [273](#)

driving factors, [266–267](#)

“fast food” mentality, [266](#)

Internet databases, [676–677](#)

Internet time, [266](#)

manageability, [267](#)

overview, [265–267](#)

versus performance, [267](#)

recoverability, [267](#)

reliability, [267](#)

response time, [266](#)

serviceability, [268](#)

Availability, ensuring

automating DBA functions, [290–291](#)

clustering technology, [292–295](#)

database architecture, [296](#)

DB2 Data Sharing, [295](#)

high-availability features, [291](#)

- load balancing, [293](#)
- nondisruptive utilities, [288–289](#)
- NoSQL, [296](#)
- online database reorganization, [288–289](#)
- performing routine maintenance, [288–289](#)
- recommended strategy, [287](#)

Availability, problems

- application problems, [279–280](#)
- authorization problems, [280](#)
- cluster failover techniques, [276](#)
- data corruption, [280–281](#)
- data replication and propagation failures, [283](#)
- DBA mistakes, [284](#), [286](#)
- DBMS software failure, [279](#)
- disk outages, [278](#)
- human error, [284](#), [286](#)
- loss of data, [282–283](#)
- loss of database objects, [281–282](#)
- loss of entire database, [277](#)
- loss of the data center, [274–275](#)
- maintenance outages, [286–287](#)
- network problems, [275](#)
- operating system failure, [279](#)
- planned outages, [286–287](#)
- recovery issues, [284](#)
- SAN failure, [278](#)
- security problems, [280](#)
- server hardware failure, [276](#)
- server performance, [283–284](#)
- standby systems, [276](#), [277](#)
- system memory failure, [276](#)
- unplanned outages, [286–287](#)

Availability, requirements

- 24/24, [270–271](#)
- across time zones, [270–271](#)
- advanced analytics, [269–270](#)
- analytics, [268–270](#)
- business intelligence, [268–270](#)
- data warehousing, [270](#)
- decision support, [268–270](#)
- five nines, [273–274](#), [292](#)
- full time, [270–271](#)
- IT complexity, [271](#)
- maintenance window, [268](#)
- MTBF (mean time between failure), [273–274](#)
- overview, [268](#)

Availability, tools. *See also* [Standby databases](#).

- AlwaysOn features, [285](#)
- Database Definition on Demand, [289–290](#)
- DB2 HADR (high-availability disaster recovery), [285](#)
- RAC (Real Application Clusters), [294](#)
- REORG, [288–289](#)

B

- Bachmann E/R method, [112](#)
- Background processes, [326](#)
- Backup. *See also* [Archiving](#); [Disaster planning](#); [Recovery](#).
 - application failure, [516–517](#)
 - concurrent access issues, [525–527](#)
 - COPY utility, [525](#)
 - data movement, [535](#)
 - data warehouse, [656–657](#)
 - database administration tasks, [26–27](#)
 - database failure types, [516–517](#)
 - database logs, [529](#)
 - database objects, [523–524](#)

- DBMS control, [524–525](#)
- differential, [521](#)
- documenting your strategy, [536](#)
- DSN1COPY utility, [552](#)
- frequency, determining, [518](#)
- full image copy, [521–523](#)
- heterogeneous database migration, [534–535](#)
- hot *versus* cold, [527](#)
- image copies, guidelines, [519–520](#)
- importance of, [515–516](#)
- incremental, [521–523](#)
- indexes, [524](#)
- instance failure, [516–517](#), [533](#), [550](#)
- media failure, [517](#), [550](#)
- object definitions, [536–537](#)
- overview, [517–520](#)
- regulatory compliance, [508](#)
- release upgrades, [534](#)
- scheduling, [531–533](#)
- SQL Server transaction logs, [530](#)
- subsystem failure, [533](#), [550](#)
- tools for, [714–715](#)
- transaction failure, [516–517](#), [550](#)

Backup, alternatives to

- disk mirroring, [556–557](#)
- exporting data, [534–535](#)
- logical backups, [534–535](#)
- redundant data, [555–556](#)
- replication, [555–556](#)
- snapshot replication, [555–556](#)
- standby databases, [277](#), [554–555](#)
- storage management software, [535–536](#), [547](#)
- symmetric replication, [555–556](#)

- unloading data, [534–535](#)
- Backup, consistency
 - creating a recovery point, [528–529](#)
 - definition, [29](#)
 - overview, [527–528](#)
 - quiesce point, [528](#)
 - QUIESCE utility, [528](#)
- Backup files, data integrity, [411](#)
- BACKUP LOG command, [530](#)
- Batch processing, [221–222](#)
- Benchmarks, [65–66](#). *See also* [Performance](#).
- Big Data movement, [55–56](#)
- Big Three DBMS vendors, [762](#)
- BIND command, [477](#)
- Binding check constraints, [424](#)
- Bitemporal support, [179](#)
- Bitmap indexes, [155–156](#)
- Bitmaps, data page layouts, [589](#)
- Blind men and an elephant, [108](#)
- Block size, optimizing database performance, [364–365](#)
- Blocks, recovering, [553](#)
- Blogs, [780–781](#)
- BMC Software, [769](#)
- Booch, Grady, [113](#)
- Bradmark Technologies, [770](#)
- Bridge architecture drivers, [673](#)
- b-tree index, [154–155](#)
- Buffer pools, [78–79](#)
- Bulk data movement, [623–625](#). *See also* [Distributed databases](#).
- Bulk-logged recovery, [340](#), [540](#)
- Bunker Hill Corporation, [770](#)
- Business intelligence
 - availability, [268–270](#)

- tool vendors, [773](#)
- tools for, [721–724](#)
- Business Intelligence Network, [783](#)
- Business logic, [664–666](#)
- Business metadata, [689](#)
- Business service interruption, risk of, [561–563](#)
- Business time, [179–180](#)
- Business-critical applications, [562](#)

C

CA Technologies, Inc., [770](#), [772](#)

Cache

- database log, [330](#)
- Internet structure, [330](#)
- memory usage, displaying, [413–414](#)
- procedure, [329–330](#), [335](#)
- program, [79](#)
- sort, [330](#)
- system performance, [328–330](#)

Cache, data

- definition, [78–79](#)
- system performance, [329–330](#), [332–335](#)

Callable routines, [192](#)

Candidate keys, indexes, [152](#)

Candle Corporation, [770](#)

Capacity planning, tools for, [313](#)

Cardett Associates, [773](#)

Cardinality, [112](#)

Careers in DBA

- DBAjobs.com, [782](#)
- demand for DBAs, [4](#)
- evaluating a job offer, [14–15](#)
- salaries, [4–6](#)

- sample job posting, [785–791](#)
- skill and knowledge requirements, [788–790](#)
- typical responsibilities, [786–788](#)
- workload, [6](#)
- Carnival blogs, [781](#)
- Cartesian products, [402](#)
- Cascading DROPs, [251](#), [701–703](#)
- Cascading REVOKEs, [462](#), [468](#)
- Cassandra, [56](#)
- Catalog query and analysis tools, [705–707](#)
- CDB Software, [770](#)
- CD/DVD storage, DBMS requirements, [77](#)
- Centralized data management model, [668](#)
- Centralized processing, [666–667](#)
- Certegy Check Services, Inc., [496](#)
- Certification
 - and job performance, [57](#)
 - overview, [56–58](#)
 - sources of, [58](#)
- Change management
 - automation, [245](#)
 - availability, [246](#)
 - checklists, [260](#)
 - coordinating databases with applications, [260–261](#)
 - DBA scripts, [262](#)
 - free space, changing, [255–256](#)
 - impact analysis, [245](#)
 - indexes, recreating, [257n](#)
 - intelligence, [245](#)
 - overview, [243–244](#)
 - perspective of the DBA, [246–247](#)
 - planning analysis, [245](#)
 - proactivity, [245](#)

- quick, efficient delivery, [246](#)
- reasons for change, [244](#)
- regulatory compliance, [261–262](#), [508](#)
- reliable, predictable processes, [246](#)
- requesting changes, [258–260](#)
- sample scenarios, [254–257](#)
- standardized procedures, [245](#)
- success factors, [244–246](#)
- tools for, [254](#), [701–703](#)

Change management, database structures

- adding columns, [255](#), [256](#). *See also* [ALTER statements](#).
- ALTER statements, [250–252](#), [252–253](#)
- cascading DROPs, [251](#)
- changing objects, [250–252](#)
- comparing structures, [257–258](#)
- CREATE statements, [250–252](#)
- creating objects, [250–252](#)
- in database change management, [250–253](#)
- DROP statements, [250–252](#)
- dropping objects, [250–252](#)
- overview, [250–252](#)
- removing objects. *See* [DROP statements](#).

Change management, types of change

- applications, [249–250](#)
- DBMS software, [248](#)
- hardware configuration, [248](#)
- logical design, [248–249](#)
- overview, [247](#)
- physical database structures, [250](#)
- physical design, [248–249](#)

Change requests, [258–260](#)

Check conditions, [420](#)

Check constraints

- benefits of, [420–421](#)
- binding, [424](#)
- check conditions, [420](#)
- constraint names, [420](#)
- definition, [28](#), [419](#)
- examples, [421–423](#)
- nulls, [423–426](#)
- versus* referential integrity, [441–442](#)
- relational nulls, [423–426](#)
- rules, [424](#)
- semantic data integrity, [419–426](#)

CHECK utility (DB2), [411](#)

CHECKALLOC option, [413](#)

CHECKCATALOG option, [413](#)

CHECKDB option, [413](#)

CHECKFILEGROUP option, [413](#)

Checkpoint/restart, tools for, [725](#)

CHECKTABLE option, [412](#)

Chen E/R method, [112](#)

Child tables, referential integrity, [433–434](#)

CKPT (checkpoint) process, [326](#)

Client computers, [665](#)

Client-based drivers, [673](#)

Client/server computing. *See also* [Database connectivity](#); [Network traffic](#).

- application servers, [664](#)
- applications, types of, [667–670](#)
- business logic, [664–666](#)
- centralized data management model, [668](#)
- centralized processing, [666–667](#)
- client computers, [665](#)
- cooperative processing, [667](#)
- database management systems, [664–666](#)
- database servers, [664](#)

- decentralized user presentation model, [668](#)
- definition, [663](#), [665](#)
- distributed data management model, [668–669](#)
- distributed processing, [666–667](#)
- distributed user presentation model, [667–668](#)
- distributing tasks across a network, [668](#)
- fat clients, [670](#)
- file servers, [664](#)
- multitier implementation, [669–670](#)
- network traffic, [670–674](#)
- performance problems, [670–674](#)
- presentation logic, [664–666](#)
- print servers, [664](#)
- recommended hardware, [666](#)
- server computers, [665–666](#)
- software layers, [664–666](#)
- thin clients, [670](#)

Cloud computing, effect on DBAs, [53–55](#)

Cloud database systems, [74–75](#)

Cluster failover techniques, [276](#)

Cluster ratios, [369](#)

Clustering. *See also* [Interleaving data](#).

- definition, [71](#), [94](#)
- indexes, [159–160](#)
- optimizing database performance, [356–358](#)
- shared-disk, [72](#)
- shared-nothing, [71–72](#)
- standards, [94](#)
- technology for availability, [292–295](#)
- types of, [71–72](#)

COBIT, [509–510](#)

CODASYL (Conference on Data Systems Languages), [754–755](#)

Codd, E. F., [128](#)

Code

- design review, [238](#)
- memory usage, displaying, [413](#)

Code generators

- creating SQL, [191–192](#)
- SQL tuning, [405](#)

Cogit, [770](#)

Cold backup, [527](#)

Column-oriented data models, [756](#)

Columns

- adding, [251](#), [255](#), [256](#)
- constraints, [144](#)
- deleting, [251](#)
- fixed-length, [144](#)
- identity property, [145](#)
- nullability, specifying, [144](#)
- ordering, [146](#)
- transforming attributes to, [142–143](#)
- unique identifiers for, [145](#)
- variable length, [144](#)

COM, SQL, [193](#)

Combined tables, optimizing database performance, [356](#)

COMMIT statements

- batch processing, [221–222](#)
- saving transaction changes, [205](#)
- SQL tuning, [404–405](#)
- two-phase, [631](#)

COMMITTED READ isolation, [216–217](#)

Communications, standards, [98](#)

Compliance, tools for, [716–721](#)

Comprehensive auditing methods, [494](#)

Compression

- data warehouse, [644](#)

- database design, [149–150](#)
- disaster planning, backup, [575](#)
- optimizing database performance, [361–362](#)
- performance tuning, [314](#)
- tools for, [726–727](#)

Computer Associates, [773](#)

Computer Associates International, [771](#)

Compuware Corporation, [770](#)

Conceptual data modeling, [125–128](#)

Conceptual design review, [233–235](#)

Concurrency, unloading data, [619](#)

Concurrency control, purpose of, [758](#)

Conference on Data Systems Languages (CODASYL), [754–755](#)

config.ora file, [325](#)

Configuring the DBMS. *See also* [Installing the DBMS](#); [Upgrading the DBMS](#).

- default parameters, [80](#)
- performance tuning. *See* [System performance](#), DBMS installation and configuration.
- system parameters, [80](#)

Confio Software, [770](#)

Connection pooling, [674](#)

Connectivity. *See* [Database connectivity](#).

Consistency, transactions, [206](#)

Consistency checking, [412–413](#)

Constraint names, [420](#)

Constraints

- check, [28](#)
- columns, [144](#)
- data integrity, [28](#)
- enforcing while loading data, [615](#)
- referential, [28](#)
- unique, [28](#)

Consultants, Web sites, [779–780](#)

Consumption sources, [330–331](#)

Contention

- performance factor, [301](#)
- performance monitoring and tuning, [23](#)
- system performance, [341–342](#)

Contingency planning. *See* [Disaster planning](#).

Control files, [325](#)

Converting data types while loading data, [616](#)

Cooperative processing, [667](#)

COPY utility, [525](#)

Copying data. *See also* [Loading data](#); [Unloading data](#).

- bulk data movement, [623–625](#)
- EXPORT utility, [622–623](#)
- IMPORT utility, [622–623](#)
- to multiple databases. *See* [Distributed databases](#).

CoSort/IRI, [773](#)

Cost-based optimization *versus* rule-based, [344](#)

Costs of

- CPU, relational optimization, [376](#)
- of data breaches, [450](#)
- I/O, relational optimization, [376](#)
- ownership, [67](#)
- performance, across the ADLC, [307](#)
- poor data quality, [488](#), [489](#)
- regulatory compliance, [485](#)
- regulatory non-compliance, [488](#)
- upgrading the DBMS, [84](#), [85](#)

CouchDB databases, [56](#), [766](#)

CPU parallelism, [391](#)

CREATE statements, [250–252](#)

CREATE TABLE statements, [436–437](#)

Creating objects. *See* [CREATE statements](#).

Critical applications, [562–563](#)
Criticality of data, ranking, [562–563](#)
Cursor, SQL, [190](#)
Cursor stability, [216–217](#)

D

DA (data administration), [15–18](#), [19](#)
DAMA (Data Management Association), Web site, [783](#)
Darwin Professional Underwriters, [450](#)
Data. *See also* [Backup](#); [Disaster planning](#); [Recovery](#).
 abstraction levels, DBMS, [757](#)
 breaches, [449–450](#)
 cleansing, [645–649](#)
 compression, [644](#)
 content, [654](#)
 corruption, [280–281](#)
 definition, [686](#)
 dictionaries, [695–696](#)
 encryption. *See* [Encryption](#).
 freshness, [654](#)
 independence, [757](#)
 latency, [574](#), [654](#)
 length, semantic data integrity, [417–418](#)
 moving. *See* [Moving data](#).
 placement, distributed databases, [629](#)
 privacy, policies and statutes, [11](#)
 profiling, [489](#), [719–720](#)
 protection, tools for, [716–721](#)
 quality, [488–489](#), [648](#). *See also* [Data integrity](#).
 rate of growth, [581](#)
 record layouts, data page layouts, [590](#)
 replication and propagation failures, [283](#)
 rows, data page layouts, [588–589](#)

security, DBMS, [758](#)

stewardship, [688](#)

usage, [655](#)

Data access

DBMS, [758–759](#)

to noncurrent versions, [177–180](#)

tracking, [490–493](#)

Data administration (DA), [15–18](#), [19](#)

Data Administration Newsletter, [782](#)

Data administration standards, [98–99](#)

Data and Technology Today blog, [781](#)

Data Control Language (DCL), [456–457](#)

Data Definition Language (DDL), [177](#), [250–252](#)

Data Dictionary. *See* [System catalog](#).

Data Encryption Standard (DES), [472](#)

Data files, [325](#)

Data governance. *See also* [Regulatory compliance](#).

IT Governance Institute, [509](#)

overview, [489–490](#)

tools for, [716–721](#)

Data Governance blog, [781](#)

Data integrity. *See also* [Data](#), [quality](#).

backup consistency, [29](#)

constraints, [28](#)

data cleansing, [646](#)

data warehouse, [646](#)

database administration tasks, [27–29](#)

DBMS, [758–759](#)

index consistency, [29](#)

pointer consistency, [29](#)

problems, loading data, [615](#)

types of, [409–410](#)

Data integrity, database structure

- backup files, [411](#)
- consistency checking, [412–413](#)
- database checking, [413](#)
- headers, [411](#)
- memory usage, [413–414](#)
- page header corruption, [411](#)
- problem management, [411–414](#)
- types of problems, [410–411](#)
- utilities for, [411–414](#)

Data integrity, semantic

- check constraints, [419–426](#)
- data length, [417–418](#)
- data types, [417–418](#)
- default values, [419](#)
- DQS (Data Quality Services), [415](#)
- entity integrity, [416–417](#)
- example, [28](#)
- overview, [414–415](#)
- primary key constraints, [416–417](#)
- triggers, [426–433](#)
- UDT (user-defined data types), [418–419](#)
- unique constraints, [417](#)
- unique entity identification, [416–417](#)

Data life cycle, [499–500](#)

Data Management Association (DAMA), Web site, [783](#)

Data mart, [638](#)

Data masking and obfuscation. *See also* [Encryption](#).

- definition, [496–497](#)
- encryption, [497](#)
- nulling out, [498](#)
- number and date variance, [497](#)
- shuffling, [497](#)
- substitution, [497](#)

table-to-table synchronization, [498](#)

techniques for, [497–498](#)

tools for, [720](#)

Data mining, [639](#)

Data modeling

concepts, [108–113](#)

conceptual, [125–128](#)

DA (data administration), [17–18](#)

DBA tasks, [33](#)

definition, [107](#)

enterprise data model, [109](#)

E/R (entity relationship diagram), [110–113](#)

homonyms, [118](#)

importance of, [107](#)

issues, [135–136](#)

logical, [125–128](#)

physical, [125–128](#)

rules for, [110](#)

synonyms, [118](#)

tool vendors, [771–772](#)

tools for, [700–701](#)

types of, [125–128](#)

Data modeling, components

attributes, [115–119](#)

entities, [113–115](#)

keys, [120–122](#)

relationships, [122–123](#)

Data models

CODASYL (Conference on Data Systems Languages), [754–755](#)

column-oriented, [756](#)

DBMS, [754–755](#), [756](#)

definition, [754](#)

denormalization, [163](#)

- hierarchical, [754–755](#)
- network, [754–755](#)
- NoSQL system, [756](#)
- object-oriented, [754–755](#)
- operations, [754](#)
- relational, [754–755](#)
- relations, [755](#)
- structure, [754](#)

Data page layouts

- allocation pages, [589](#)
- allocation units, [589](#)
- bitmaps, [589](#)
- data record layouts, [590](#)
- data rows, [588–589](#)
- header information, [592](#)
- index key values, [592](#)
- index page layouts, [592–594](#)
- offset and adjust tables, [592](#)
- offset tables, [588–590](#)
- overview, [588–589](#)
- page header, [588–589](#)
- page pointer, [592](#)
- row data, [590](#)
- row header, [590](#)
- row length, [592](#)
- sample, [589](#)
- space page map, [589](#)
- table size, calculating, [591–592](#)
- transaction logs, [594–595](#)

Data resource management (DRM), [40–42](#)

Data retention

- DBA source materials, rule of thumb, [736–737](#)
- disaster planning backup, [571](#)

Data sets. See [Files and data sets](#).

Data spaces, database design, [148](#)

Data types

attributes, [116](#)

converting, while loading data, [616](#)

semantic data integrity, [417–418](#)

Data warehouse

administrators, [36–37](#)

analytical *versus* transaction processing, [638–640](#)

availability, [270](#)

data mining, [639](#)

definition, [637–638](#)

design, [641–644](#)

dimensions, [639](#)

DSS (decision support systems), [639](#)

facts, [639](#)

Information Center, [640](#)

metadata, [688](#)

OLAP (online analytical processing), [639](#)

OLAP *versus* OLTP, [640](#)

tools for, [721–724](#)

Data warehouse, administering

backup and recovery, [656–657](#)

data cleansing, [645–649](#)

data compression, [644](#)

data content, [654](#)

data freshness, [654](#)

data integrity problems, [646](#)

data latency, [654](#)

data movement, [644–645](#)

data quality issues, [648](#)

data usage, [655](#)

data warehouse design, [641–644](#)

- denormalization, [643](#)
- financial chargeback, [655–656](#)
- focus on technology, [641](#)
- identifying unused data, [655](#)
- meeting business requirements, [657](#)
- metadata, [654](#)
- operational problems, [648–649](#)
- overview, [640–641](#)
- purging data, [655](#)
- scalability, [649](#)
- size issues, [649](#)
- snowflake schema, [643](#)
- standardizing default values, [647](#)
- star schema, [641–643](#)

Data warehouse, performance

- automatic summary tables, [652–653](#)
- data management, [650](#)
- extract performance, [650](#)
- indexes, [651](#)
- materialized query tables, [652–653](#)
- materialized views, [653](#)
- monitoring, [652](#)
- perspectives on, [650](#)
- query performance, [650](#)
- server performance, [650](#)

The Data Warehousing Information Center, [783](#)

The Data Warehousing Institute, [783](#)

Database administration. *See also* [DBA \(database administrator\)](#).

- importance of, [3–4](#)
- management discipline, [9–14](#)

Database administration tasks. *See also specific tasks.*

- availability, [24](#)
- backup and recovery, [26–27](#)

- data integrity, [27–29](#)
- database design, [21–22](#)
- governance and regulatory compliance, [26](#)
- jack-of-all-trades, [29–31](#)
- performance monitoring and tuning, [22–23](#)
- security and authorization, [24–25](#)

Database administrator (DBA). *See* [DBA \(database administrator\)](#).

Database applications, designing. *See also* [Design review](#); [SQL \(Structured Query Language\)](#); [Transactions](#).

- ADO.NET, [194–195](#)
- application infrastructure, [193–194](#)
- hardware environment, [193–194](#)
- issues, [186](#)
- J2EE (Java 2 Enterprise Edition), [195–196](#), [198](#)
- Java program types, [196–197](#)
- .NET framework, [194–195](#), [198](#)
- overview, [185–186](#)
- Ruby on Rails, [198](#)
- software environment, [193–194](#)

Database architects, DBAs as, [32–33](#)

Database connectivity. *See also* [Client/server computing](#); [Internet](#); [Network traffic](#); [Web services](#).

- application servers, [664](#)
- business issues, [662](#)
- client/server computing, [663–666](#)
- database servers, [664](#)
- downsizing, [662](#)
- file servers, [664](#)
- history of, [661–662](#)
- print servers, [664](#)
- rightsizing, [662](#)
- upsizing, [662](#)

Database Definition on Demand, [289–290](#)

Database design. *See also* [Indexes](#); [Views](#).

compression, [149–150](#)

data spaces, [148](#)

database administration tasks, [21–22](#)

domains, transforming to data types, [143–144](#)

for e-business, [677–680](#)

entities, transforming to tables, [142](#)

filegroups, [149](#)

logical model to physical database, [141–150](#)

overview, [141–142](#)

physical data structure, [147–150](#)

primary keys, [144](#)

raw files, [149](#)

referential constraints, [146–147](#)

referential integrity, [146–147](#)

row size, specifying, [148](#)

storage requirements, [148](#)

tablespaces, [148](#)

temporal requirements, [177–180](#)

Database design, columns

constraints, [144](#)

fixed-length, [144](#)

identity property, [145](#)

nullability, specifying, [144](#)

ordering, [146](#)

transforming attributes to, [142–143](#)

unique identifiers for, [145](#)

variable length, [144](#)

Database drivers, [672–674](#)

Database environments. *See also* *specific environments*.

education, [101](#)

integration testing, [101](#)

multiplatform issues, [42–43](#)

- production *versus* test, [44–46](#)
- quality assurance testing, [101](#)
- unit testing, [101](#)
- user acceptance testing, [101](#)

Database files, Oracle, [325](#)

Database gateways, [671–672](#)

Database ID. *See* [Users](#), names.

Database logs

- active, [529](#)
- archiving, [77–78](#), [339](#), [529](#)
- backing up, [529](#), [530](#)
- bulk-logged recovery, [340](#)
- configuring, [338–339](#), [340](#)
- DBMS, [758](#)
- definition, [336](#)
- disabling, [341](#)
- disabling while loading data, [617](#)
- disaster planning, backup, [570–571](#)
- filling up, [339](#)
- full recovery, [340](#)
- log archival process, [529](#)
- log off-loading, [339](#)
- log-based auditing, [493–495](#)
- “out of space” conditions, [339–341](#)
- placement for optimizing performance, [363](#)
- during recovery, [338](#), [340](#)
- recovery models, [340](#)
- selecting candidates for, [339–341](#)
- simple recovery, [340](#)
- system checkpoints, [337](#)
- system performance, [336–341](#)
- transaction logs, [336](#)
- types of information on, [337](#)

write-ahead, [337](#)

Database management system (DBMS). *See* [DBMS \(database management system\)](#).

Database performance. *See also* [Optimizing database performance](#); [Reorganizing databases](#); [SQL tuning](#).

80/20 (Pareto) rule, [302](#)

versus availability, [267](#)

common problems, [302–304](#)

contention, [301](#)

cost, across the ADLC, [307](#)

definition, [23](#), [300–302](#)

diagnosing, [302–304](#)

estimating, [307–308](#)

guidelines, [315–316](#)

historical trends, [308](#)

main factors, [301–302](#)

overview, [299–302](#)

resources, [301](#)

SLM (service-level management), [308–311](#)

throughput, [301](#)

tools for, [711](#)

tracker tables, [308](#)

tuning SQL, [303–304](#)

workload, [301](#)

Database performance, managing

analysis, [305](#)

components of, [304–306](#)

definition, [304–306](#)

versus monitoring, [304–306](#)

overview, [304–306](#)

reactive *versus* proactive, [306](#)

Database performance, monitoring

contention, [23](#)

- database administration tasks, [22–23](#)
- factors affecting, [22–23](#)
- resources, [22](#)
- throughput, [22](#)
- tools for, [313](#)
- workload, [22](#)

Database performance, tuning

- application, [312](#)
- caching, [314](#)
- capacity planning, [313](#)
- compression, [314](#)
- contention, [23](#)
- database, [312](#)
- database administration tasks, [22–23](#)
- estimation, [313](#)
- factors affecting, [22–23](#)
- monitoring, [313](#)
- reorganizing databases, [314](#)
- resources, [22](#)
- sorting, [314](#)
- SQL analysis and tuning, [313](#)
- system, [311](#)
- throughput, [22](#)
- tools for, [313–315](#)
- workload, [22](#)

Database servers

- definition, [664](#)
- hosting, [675](#)
- location, upgrading, [88](#)

Database Site portal, [781](#)

Database Trends and Applications, [779](#)

Database views. *See* [Views](#).

Database wire drivers, [674](#)

Database writer (DBWR) process, [326](#)
Database-coupled application logic, [46–50](#)

Databases

- architecture for availability, [296](#)
- change management. *See* [Change management](#).
- checking for data integrity, [413](#)
- comparison, tools for, [703–704](#)
- DBA staffing requirements, [37](#)
- versus* DBMS, [7–8](#)
- definition, [7](#), [753–754](#)
- dropping, [250–252](#)
- links, [94](#)
- logic, managing, [46](#)
- maintenance, authorization, [467](#)
- management systems, [664–666](#)
- object privileges, [459](#)
- objects, backing up, [523–524](#)
- Oracle, [325](#)
- structures, comparing, [257–258](#)
- tools for. *See* [Tools](#).
- users, security, [455–456](#)

DataBee, [771](#)

Datanamic, [771](#)

DB2 (IBM)

- blogs, [780](#)
- IDUG (International DB2 User Group), [740](#), [783](#)
- nonstandard database objects, [94](#)
- vendor contact, [63](#)
- Web site, [778](#)

DB2 Catalog. *See* [System catalog](#).

DB2 Data Sharing, [295](#)

DB2 EDM pools, [335](#)

DB2 HADR (high-availability disaster recovery), [285](#)

DBA (database administrator). *See also* [Database administration](#).

authorization, [467](#)

versus DA, [15–18](#), [19](#), [21](#)

a day in the life of, [12–14](#)

demand for, [4](#)

description, [1–3](#)

job scope, defining, [42–43](#)

jobs. *See* [Careers in DBA](#).

multiplatform issues, [42–43](#)

reporting structures, [40–42](#)

responsibilities, [12](#), [786–788](#)

versus SA, [21](#)

skill and knowledge requirements, [788–790](#)

staffing, [37–40](#)

standards and procedures, [98–99](#)

tools for. *See* [Tools](#), for DBAs.

typical responsibilities, [786–788](#)

workload, [6](#), [12–14](#)

DBA (database administrator), rules of thumb

accessibility to coworkers, [745](#)

analysis, [741–742](#)

automation, [737–739](#)

being prepared, [743](#)

calm in the face of adversity, [742–743](#)

documenting your work, [735–736](#)

effective use of resources, [745–746](#)

focus, [741–742](#)

investing in professional advancement, [747–748](#)

retaining source materials, [736–737](#)

sharing knowledge, [739–741](#)

simplification, [741–742](#)

technical education, [746–747](#)

Twitter, as a resource, [741](#)

- understanding your business, [743–745](#)
- user group associations, [740](#)

DBA (database administrator), types of

- application, [34–35](#)
- data modeler, [33](#)
- data warehouse administrator, [36–37](#)
- database analysts, [33](#)
- database architects, [32–33](#)
- performance analysts, [36](#)
- system, [31–32](#)
- task-oriented, [36](#)
- technical focus *versus* business, [31–32](#)

DBA Direct, [779](#)

DBAjobs.com, [782](#)

dBase, [767](#)

DBCC utility, options, [412–414](#)

DBE Software, [771](#)

DBI Software, [771](#)

dbMaestro, [771](#)

DBMS (database management system)

- architectures, [68–71](#)
- atomicity, [758](#)
- availability problems, [279](#)
- buying. *See* [Vendors](#), [DBMS](#).
- clustering. *See* [Clustering](#).
- concurrency control, [758](#). *See also* [Locking](#).
- data abstraction levels, [757](#)
- data access, [758–759](#)
- data independence, [757](#)
- data integrity, [758–759](#)
- data models, [754–755](#), [756](#)
- data security, [758](#)
- versus* database, [7–8](#)

- database logging, [758](#)
- definition, [8](#), [753–754](#)
- departmental architecture, [70](#)
- durability, [758](#)
- enterprise architecture, [69](#)
- mobile architecture, [70](#)
- organizational strategy. *See* [Strategies for DBMS](#).
- personal architecture, [70](#)
- proliferation, [73](#)
- upgrading, [87–88](#)
- vendors. *See* [Vendors, DBMS](#).

DBWR (database writer) process, [326](#)

DCL (Data Control Language), [456–457](#)

DDL (Data Definition Language), [177](#), [250–252](#)

Deadlock detection, [341](#)

Deadlocks, [214–215](#), [342](#)

Debugging, tools for, [726](#)

Decentralized user presentation model, [668](#)

Decision support, availability, [268–270](#)

Decision support systems (DSS), [639](#)

Defragmenting indexes, [413](#)

DELETE privileges, [458–459](#)

DELETE rule, [435–436](#)

DELETE statements

- modifying temporal data, [180](#)
- in triggers, [429](#)

DELETE trigger, [438–441](#)

Deleting objects. *See also* [DROP statements](#).

- columns, [251](#)
- purging data, data warehouses, [655](#)
- purging databases, *versus* archiving, [501](#)
- rows, [435–436](#)

Denormalization

- benefits, evaluating, [175](#)
- combined tables, [168](#)
- data warehouse, [643](#)
- derivable data, [170–171](#)
- description, [160–161](#)
- evaluating the need for, [161–162](#), [174–175](#)
- hierarchies, [171–173](#)
- identifying candidates for, [162–163](#)
- Internet databases, [680](#)
- issues, [161–162](#)
- logical data models, [163](#)
- mirror tables, [165](#)
- optimizing database performance, [355–356](#)
- overview, [161–163](#)
- physical implementation requirements, [173](#)
- prejoined tables, [164](#)
- redundant data, [168–169](#)
- repeating groups, [169–170](#)
- report tables, [164–165](#)
- speed tables, [172–173](#)
- split tables, [165–166](#)
- splitting text columns, [166–168](#)
- types of, [174](#)

Density, relational optimization, [377](#)

Departmental DBMS architectures, [70](#)

Deprecated features, [85n.7](#)

Derivable data, [170–171](#), [356](#)

Derived data, storing *versus* calculating, [170–171](#)

DES (Data Encryption Standard), [472](#)

Design review

- guidelines, [102](#), [228–229](#)
- output, [239–240](#)
- overview, [227–228](#)

purpose of, [228](#)

Design review, participants

knowledge and skills required, [232](#)

leader of, [229–230](#)

mediator, [230–231](#)

mentorship and knowledge transfer, [240–241](#)

recommended personnel, [231](#)

remote staff, [232](#)

scribe, [230](#)

Design review, types of

in the ADLC, [234](#)

code, [238](#)

conceptual, [233–235](#)

logical, [235](#)

organizational, [237](#)

overview, [233](#)

physical, [236](#)

post-implementation, [239](#)

pre-implementation, [239](#)

Designing

applications. *See* [Database applications, designing](#).

databases. *See* [Database design](#).

Determinant, [135](#)

Devices, naming conventions, [364](#)

Diagramming entity relationships. *See* [E/R \(entity relationship diagram\)](#).

Differential backup, [521](#)

Dimensions, data warehouse, [639](#)

Direct index lookup, [383](#)

Dirty read, [216–217](#)

DISABLE option, [477](#)

Disabling

database logs, [341](#)

passwords, [453](#)

Disaster, definition, [559–560](#)

Disaster planning. *See also* [Backup](#); [Recovery](#).

business-critical applications, [562](#)

critical applications, [562–563](#)

criticality of data, ranking, [562–563](#)

lengthy outages, [568](#)

need for, [559–563](#)

noncritical applications, [563](#)

prevention, [575–576](#)

required applications, [563](#)

very critical applications, [562](#)

Web sites about, [576](#)

Disaster planning, backup

compression, [575](#)

data latency, [574](#)

data retention, [571](#)

database logs, [570–571](#)

encryption, [575](#)

important files and data, [574–575](#)

indexes, [570](#)

order of recovery, [574](#)

over a WAN (wide-area network), [573](#)

post-recovery image copies, [575](#)

remote mirroring, [573](#)

standby databases, [573](#)

storage management software, [572–573](#)

on tape, [570–571](#)

Disaster planning, recovery

off-site locations, [564](#)

personnel, [569](#)

plan content, [566](#)

recovery site, choosing, [564](#)

rehearsing, [567–569](#)

- team members, [569](#)
- testing your plan, [567–569](#), [574](#)
- written plans, [564–566](#)

Disaster planning, risk

- assessing, [561–563](#)
- business service interruption, [561–563](#)
- categories of, [561](#)
- financial loss, [561–563](#)
- legal responsibilities, [561–563](#)

Discovering attributes and entities, [124–125](#)

Disk drives

- MTBF (mean time between failures), [580](#)
- overview, [580](#)

Disk storage

- DBMS requirements, [76–78](#)
- SSDs (solid state devices) *versus* traditional disks, [323–324](#)
- system performance, [322–324](#)

Disks. *See also* [Storage management](#).

- allocation, optimizing database performance, [364](#)
- fragmentation, [595](#)
- JBOD (just a bunch of disks), [604](#)
- mirroring, as a backup/recovery alternative, [556–557](#)
- outages, [278](#)
- performance improvement, [584–585](#)
- raw partitions *versus* file systems, [586–587](#)
- SCSI (small computer system interface), [605](#)
- short-stroking, [584–585](#)
- size terminology, [582](#)
- storage management option, [596](#)
- striping, [597](#)
- usage spikes, [580](#)

DISTINCT clause, [387](#)

Distributed data

- accessing, [630–631](#)
- distributed request, [631](#)
- distributed unit of work, [631](#)
- DRDA (Distributed Relational Database Architecture), [629–630](#)
- placement for optimum performance, [363–364](#)
- RDA (Remote Database Access), [629–630](#)
- remote requests, [630–631](#)
- remote unit of work, [630–631](#)
- standards, [629–630](#)
- two-phase COMMIT, [631](#)

Distributed data management model, [668–669](#)

Distributed databases. *See also* [Bulk data movement](#); [Copying data](#); [Moving data](#).

- autonomy, [626](#)
- characteristics of, [626](#)
- data placement, [629](#)
- definition, [626](#)
- environment, setting up, [627–629](#)
- federated multidatabase schemes, [627](#)
- isolation, [626](#)
- performance problems, [632–633](#)
- system performance, [344](#)
- transparency, [626](#)
- unfederated multidatabase schemes, [627](#)
- usage guidelines, [629](#)

Distributed processing, [666–667](#)

Distributed request, [631](#)

Distributed unit of work, [631](#)

Distributed user presentation model, [667–668](#)

Distributing tasks across a network, [668](#)

Document Type Definition (DTD), [204](#)

Documentation

- DBA activities, [735–736](#)

online standards manuals, [727–728](#)

Domains

attributes, [116](#)

transforming to data types, [143–144](#)

Downsizing, and database connectivity, [662](#)

Downtime. *See also* [Availability](#).

versus availability, [273](#)

cost of, [271–273](#)

DBA staffing requirements, [38](#)

negative publicity, [272](#)

DQS (Data Quality Services) (Microsoft), [415](#)

DRDA (Distributed Relational Database Architecture), [629–630](#)

Drivers

JDBC, [673–674](#)

ODBC, [192](#), [673](#)

DRM (data resource management), [40–42](#)

DROP statements

cascading DROPs, [251](#), [701–703](#)

in database change management, [250–252](#)

Dropped database objects, recovering, [552–553](#)

Dropping

database objects, [250–252](#)

tables, [250–252](#)

DSN1COPY utility, [552](#)

DSNZPARM parameter, [80n.4](#)

DSS (decision support systems), [639](#)

DTD (Document Type Definition), [204](#)

Duplicate values, relational optimization, [377](#)

Durability

DBMS, [758](#)

transactions, [206](#)

Dynamic SQL, [201](#)

E

Ebbers, Bernard, [485](#)

E-business. *See also* [Internet](#).
 effects on DBAs, [50–51](#)
 infrastructure, [52](#)

E-discovery, effects on DBA, [506–507](#)

EDM pools, [335](#)

Education
 database environment, [101](#)
 recommended courses, [103](#)

E-Government Act, [484–485](#)

80/20 (Pareto) rule, [302](#)

Elephant and blind men, [108](#)

Embarcadero Technologies, [770–771](#)

Embedded SQL, [191–192](#), [201](#)

ENABLE option, [477](#)

Encoding scheme, specifying, [620](#)

Encryption. *See also* [Data masking and obfuscation](#).
 data at rest, [472](#)
 data in transit, [472](#)
 data masking and obfuscation, [497](#)
 disaster planning, backup, [575](#)
 overview, [470](#), [472](#)
 techniques for, [472](#)
 transparent, [473](#)
 wallets, [473](#)

End-to-end performance, tools for, [713–714](#)

Enron Corporation, [485](#)

Enterprise data model, [109](#)

Enterprise DBMS architectures, [69](#)

Entities. *See also* [Relationships](#).
 adjectives as, [115](#)
 associative, [127](#)
 definition, [113](#)

- discovering, [124–125](#)
- instances, [115](#)
- naming conventions, [113](#)
- nouns as, [115](#), [124](#)
- transforming to tables, [142](#)

Entity integrity, [416–417](#)

Entity occurrences, [115](#)

Environments, system. See [Database environments](#).

Epsilon, data breach, [449](#)

E/R (entity relationship diagram)

- Bachmann method, [112](#)
- cardinality, [112](#)
- Chen method, [112](#)
- definition, [110](#)
- diagramming methods, [111–113](#)
- example, *111*
- Information Engineering method, [112](#)
- Martin method, [112](#)
- Ross method, [112](#)
- Rumbaugh method, [112](#)
- UML (Unified Modeling Language), [113](#)

Error correction coding, [599](#)

Estimating

- memory requirements, [331–332](#)
- performance, [307–308](#)
- tools for, [313](#)

ETL (extract, transfer, load), [623–625](#), [721–723](#)

ETL tool vendors, [773](#)

E-vailability, [676–677](#)

EXCEPT clause, [388](#)

Exclusive locks, [213](#)

EXECUTE privileges, [460](#)

EXPLAIN command, [394–398](#), [712](#)

Exploiting *versus* supporting, [91](#)
EXPORT utility, [622–623](#). *See also* [UNLOAD utility](#).

Exporting data

 backup/recovery alternative, [534–535](#)

 EXPORT utility, [622–623](#)

EXtensible Markup Language (XML), [204](#)

External security, [478–480](#)

Extract performance, [650](#)

F

Fabian Pascal's site, [779](#)

Facts, data warehouse, [639](#)

Fallback planning, [92](#)

“Fast food” mentality, [266](#)

Fat clients, [670](#)

Fault tolerance, [601–602](#)

Federal Rules of Civil Procedure (FRCP), [506–507](#)

Federated multidatabase schemes, [627](#)

Fiber channel, storage management option, [605](#)

File extents, reorganizing databases, [366](#)

File servers, database connectivity, [664](#)

File systems *versus* raw partitions, [586–587](#)

Filegroups

 database design, [149](#)

 definition, [94](#)

 standards, [94](#)

FileMaker, [767](#)

Files and data sets. *See also* [Storage management](#).

 optimal placement, [584–586](#)

 overview, [583–584](#)

 placement and allocation, optimizing database performance, [362–364](#)

 temporary database files, [587](#)

Fill factor. *See* [Free space](#).

Financial chargeback, [655–656](#)
Financial loss, risk of, [561–563](#)
Financial Modernization Act of 1999, [484–485](#)
Firing triggers, [428–429](#)
FISMA (Federal Information Security Management Act), [485](#)
Five-nines availability, [273–274](#), [292](#)
Fixed-length columns, [144](#)
Fixpacks and maintenance, [480–481](#)
Floating-point data
 loading, [616](#)
 unloading data, [620](#)
Focus, DBA rule of thumb, [741–742](#)
FORCEPLAN option, [398–399](#)
Forcing access path choices, [398–399](#)
Foreign key perspective, [434–435](#)
Foreign key values, [434–436](#)
Foreign keys, indexes, [151](#)
Forrester Research, [450](#)
Fragmentation
 disks, [595](#)
 indexes, [595](#)
 reorganizing databases, [366](#)
FRCP (Federal Rules of Civil Procedure), [506–507](#)
Free space
 changing, [255–256](#)
 optimizing database performance, [360–361](#)
FREEPAGE parameter, [360](#)
Full image copy, [521–523](#)
Full recovery, [340](#), [540](#)
Full-time availability, [270–271](#)

G

Gerstner, Lou, [581](#)

Giga Research Group, [582](#)
GLB (Gramm-Leach-Bliley) Act, [484–485](#), [491](#)
Governance. See [Data governance](#).
Grandite, [772](#)
GRANT statements, [456–457](#)
Granularity
 locks, [210–211](#), [219–220](#)
 triggers, [431–432](#)
GROUP BY clause, [388](#)
Groups, authorization, [468](#)
“Guilty until proven innocent,” [13](#)

H

Hadoop databases, [766](#)
Hardware
 configuration for system performance, [322–324](#)
 environment, designing, [193–194](#)
 issues, strategies for, [73–74](#)
 requirements, installing the DBMS, [76](#)
Hash function, [389](#), [390](#)
Hash joins, [379](#)
Hashed access, [389–390](#)
Hashing, randomizing, [158](#)
HBase databases, [56](#), [766](#)
Header information, data page layouts, [592](#)
Headers, data integrity, [411](#)
Health Net Federal Services, [450](#)
Heterogeneous database migration, [534–535](#)
Hibernate, ORM library, [200](#)
Hierarchical data models, [754–755](#)
Hierarchies, [171–173](#)
HIPAA (Health Insurance Portability and Accountability Act), [484–485](#), [491](#)
Hit Software, [771](#)

Homonyms, in data modeling, [118](#)
Horizontal restriction, [469](#)
Hostile databases, [678–679](#)
Hosting database servers, [675](#)
Hot backup, [527](#)
Human error, availability problems, [284](#), [286](#)
100 Year Archive Requirements Survey, [503n](#)
Hybrid joins, [379](#)

I

IBM Corporation
 DB2 Web site, [778](#)
 DBMS vendor, 63–64, [762](#), [765](#)
 rate of data growth, [581](#)
 tool vendor, [772](#), [773](#)
IBM Data Management, [779](#)
IDC Corporation, [581](#)
IDE (integrated development environment), [191–192](#)
Identity property, columns, [145](#)
Identity values, system performance, [344](#)
Idera, [771](#)
IDMS (Cullinet), [765](#)
IDUG (International DB2 User Group), [740](#), [783](#)
IIUG (International Informix Users Group), [740](#), [783](#)
ILM (information life cycle management), [606](#)
Image copies
 backup guidelines, [519–520](#)
 backups, unloading data from, [619](#)
 disaster planning, post-recovery, [575](#)
Impact analysis, change management, [245](#)
IMPORT utility, [622–623](#). *See also* [LOAD utility](#).
Importing data, [622–623](#)
IMS (IBM), [765](#)

Incremental backup, [521–523](#)

Independent Oracle Users Group (IOUG), [740](#)

Index covering, [386–387](#)

Indexed access, [382–389](#)

Indexes

- absence of, [151](#)

- avoiding, [354](#)

- avoiding sorts, [387–388](#)

- backup, [524](#)

- based on workload, [152](#)

- bitmap, [155–156](#)

- b-tree, [154–155](#). *See also* [Partitioned index](#); [Reverse key index](#).

- candidate keys, [152](#)

- clustering, [159–160](#)

- consistency, [29](#)

- costs of, [153–154](#)

- data warehouse, [651](#)

- defragmenting, [413](#)

- designing, [150–154](#)

- disaster planning, backup, [570](#)

- file placement, [584](#)

- foreign keys, [151](#)

- fragmentation, [595](#)

- indexing by object, [152](#)

- index-only access, [152](#)

- key values, [592](#)

- leaf pages, [155](#)

- locking, [212](#)

- nodes, [155](#)

- optimal number per table, [353](#)

- optimizing database performance, [352–355](#)

- ordered, [157](#)

- overloading, [355](#)

- page layouts, [592–594](#)
- partitioned, [157](#). *See also* [b-tree index](#).
- primary keys, [151](#)
- recovery, [550–551](#)
- recreating, [257n](#)
- reorganizing, [369–370](#)
- reverse key, [156–157](#). *See also* [b-tree index](#).
- screening, [386](#)
- size, calculating, [592–594](#)
- sorting, [152](#)
- table scans, [151](#)
- unused, dropping, [153](#)

Indexing by object, [152](#)

Index-only access, [152](#), [386–387](#)

Informatica, [773](#)

Information, definition, [687](#)

Information Builders, [773](#)

Information Center, [640](#)

Information Engineering, E/R method, [112](#)

Information life cycle management (ILM), [606](#)

Information Management, [779](#)

Information Schema. *See* [System catalog](#).

Informix, [763](#)

- IIUG (International Informix Users Group), [740](#), [783](#)
- vendor contact, [64](#)
- Web site, [778](#)

InfoTel Corporation, [771](#)

Ingres, [763–764](#)

INIT.ORA, [80n.4](#)

Inner table, [379](#)

INSERT privileges, [458–459](#)

INSERT rule, [434–436](#)

INSERT statements

- modifying temporal data, [180](#)
- recording in the transaction log, [341](#)
- in triggers, [429](#)

INSERT trigger, [438–441](#)

Installing the DBMS. *See also* [Configuring the DBMS](#); [Upgrading the DBMS](#).

- connecting to infrastructure software, [81](#)
- hardware requirements, [76](#)
- memory requirements, [78–79](#)
- in multiple environments, [82](#)
- performance tuning. *See* [System performance](#), DBMS installation and configuration.
- prerequisites, [75–76](#)
- storage requirements, [76–78](#)
- verifying the install, [81](#)

Instance failure, backup, [516–517](#), [533](#), [550](#)

Instances

- entities, [115](#)
- Oracle databases, [325](#)

INSTEAD OF trigger, [432](#)

Integrated development environment (IDE), [191–192](#)

Integrated metadata sources, [690](#)

Integration testing, [101](#)

Integrity. *See* [Data integrity](#); [RI \(referential integrity\)](#).

Intelligence, change management, [245](#)

Intent locks, [213](#)

Interconnected databases, [676–680](#)

Interleaving data. *See also* [Clustering](#).

- optimizing database performance, [360](#)
- performance design, [160](#)

International DB2 User Group (IDUG), [740](#), [783](#)

International Informix Users Group (IIUG), [740](#), [783](#)

International issues, Internet databases, [679](#)

International Oracle Users Group, [783](#)

International Sybase User Group (ISUG), [740](#), [783](#)

Internet, and e-business

effect on DBAs, [50–52](#)

infrastructure, [52](#)

Internet, database connectivity. *See also* [Web services](#).

availability, [676–677](#)

denormalization, [680](#)

designing for e-business, [677–680](#)

effect on DBA duties, [676–680](#)

e-availability, [676–677](#)

hostile databases, [678–679](#)

hosting database servers, [675](#)

interconnected databases, [676–680](#)

international issues, [679](#)

Internet time, [677](#)

key design, [679](#)

normalization, [679](#)

RAD (rapid application development), [677](#)

Internet resources

blogs, [780–781](#)

database portals, [781–782](#)

industry standards, [782](#)

jobs, [782](#)

mailing lists, [776–778](#)

Usenet newsgroups, [775–776](#)

Internet resources, Web sites

consultants, [779–780](#)

magazines, [778–779](#)

Mullins, Craig, [780](#)

user group associations, [740](#)

vendors, [778](#)

Internet time, [266](#), [677](#)

INTERSECT clause, [388](#)

Invasive performance tools, [710](#)

I/O parallelism, [390–391](#)

IOUG (Independent Oracle Users Group), [740](#)

ISO Web site, [782](#)

Isolation

 distributed databases, [626](#)

 levels, [216–218](#)

 transactions, [206](#)

ISQL, [81n.5](#)

ISUG (International Sybase User Group), [740](#), [783](#)

IT complexity, and availability, [271](#)

IT Governance Institute, [509](#)

J

J2EE (Java 2 Enterprise Edition), [195–196](#), [198](#)

Jack-of-all-trades, [29–31](#)

Jacobson, Ivar, [113](#)

Java

applets, [196–197](#)

applications, [196–197](#)

choosing a program type, [196–197](#)

Hibernate, ORM library, [200](#)

LINQ (Language Integration Query), [200](#)

NHibernate, ORM library, [200](#)

program types, [196–197](#)

servlets, [196–197](#)

JBOD (just a bunch of disks), [604](#)

JDBC (Java Database Connectivity)

bridge architecture drivers, [673](#)

client-based drivers, [673](#)

database wire drivers, [674](#)

drivers, [673–674](#)

network protocol architecture drivers, [673–674](#)

overview, [192–193](#)

Pure Java drivers, [673–674](#)

Type 1 drivers, [673](#)

Type 2 drivers, [673](#)

Type 3 drivers, [673–674](#)

Type 4 drivers, [673](#)

Job scheduling, security, [479](#)

Jobs in DBA. See [Careers in DBA](#).

Join order, [381](#)

Joining tables

hash join, [379](#)

hybrid join, [379](#)

- inner table, [379](#)
- join order, [381](#)
- merge-scan join, [379–380](#)
- nested-loop join, [379–380](#)
- outer table, [379](#)
- qualifying rows, [380](#)
- relational optimization, [379–381](#)

Joins, SQL, [189](#)

K

Kernel memory usage, displaying, [413](#)

Keys

- candidate, [115](#), [121](#)
- description, [120](#)
- designing for Internet connectivity, [679](#)
- foreign, [121–122](#)
- primary, [121](#)

Knowledge, definition, [687](#)

Kozlowski, Dennis, [485](#)

L

Lay, Ken, [485](#)

LBAC (label-based access control), [463–465](#)

Leader, design review, [229–230](#)

Leaf distance, reorganizing databases, [370](#)

Leaf pages, [155](#)

Legal responsibilities, risks, [561–563](#)

LGWR (log writer) process, [326](#)

Life cycle

- ADLC (application development life cycle), [9–10](#), [12](#)
- data, [499–500](#)

LIKE logical operator, [403–404](#)

LIMIT parameter, [620–621](#)

LINQ (Language Integration Query), [200](#)

List servers, [777](#)

Listservs. *See* [Mailing lists](#).

Load balancing, [293](#)

LOAD parameters, [620](#)

LOAD utility, [614–618](#), [621–622](#). *See also* [IMPORT utility](#).

Loading data. *See also* [Unloading data](#).

- for application test beds, [621–622](#)

- converting data types, [616](#)

- data integrity problems, [615](#)

- describing the input file, [615](#)

- disabling logging, [617](#)

- efficiency, [617–618](#)

- enforcing constraints, [615](#)

- firing triggers, [615](#)

- floating-point data, [616](#)

- LOAD utility, [614–618](#), [621–622](#)

- nulls, [616](#)

Lock duration

- acquire/release specification, [218](#)

- definition, [215](#)

- isolation level, [216–218](#)

- SKIP LOCKED DATA parameter, [219](#)

- skipping locked rows, [219](#)

Lock escalation, [219–220](#)

Lock suspensions, [341](#)

Locking

- COMMITTED READ isolation, [216–217](#)

- cursor stability, [216–217](#)

- deadlocks, [214–215](#)

- description, [210](#)

- dirty read, [216–217](#)

- exclusive locks, [213](#)

- granularity, [210–211](#), [219–220](#)
- index entries, [212](#)
- intent locks, [213](#)
- levels. See [Granularity](#).
- minimizing problems, [220](#)
- passwords, [453](#)
- phantoms, [218](#)
- read locks, [212–213](#)
- REPEATABLE READ isolation, [217](#)
- SERIALIZABLE isolation, [218](#)
- shared locks, [212–213](#)
- system performance, [341–342](#)
- time outs, [213–214](#)
- types of locks, [212–213](#)
- UNCOMMITTED READ isolation, [216–217](#)
- update locks, [213](#)
- write locks, [212–213](#)

Log writer (LGWR) process, [326](#)

Logging. See [Database logs](#).

Logical backups, [534–535](#)

Logical data independence, [757](#)

Logical data modeling, [125–128](#)

Logical design review, [235](#)

Logical models, converting to a physical database, [141–150](#)

Logic-oriented security, [470](#)

Logins

- administration, rules of thumb, [453](#), [455](#)

- definition, [455–456](#)

- limiting, [455](#)

- required information, [452–453](#)

Loss of

- data, [282–283](#)

- data center, [274–275](#)

database objects, [281–282](#)

entire database, [277](#)

Lotus, [767](#)

LPARS (Logical PARtitions), [295n](#)

M

Magazines, Web sites, [778–779](#)

Mailing lists, [776–778](#)

Main-memory database management systems (MMDBMSs), [596](#)

Maintenance

outages, availability problems, [286–287](#)

patches for Oracle databases, [480–481](#)

scheduling for availability, [288–289](#)

window, availability, [268](#)

Manageability, availability, [267](#)

Mapping physical files to database tables, [350–352](#)

Martin E/R method, [112](#)

Masking data. *See* [Data masking and obfuscation](#).

Master Data Services. *See* [SQL Server](#).

Matching index scan, [383–385](#)

Materialized query tables, [652–653](#)

Materialized views, [653](#)

Mean time between failures (MTBF)

availability, [273–274](#)

disk drives, [580](#)

Media failure, backup, [517](#), [550](#)

Mediator, design review, [230–231](#)

Memory, requirements

buffer pools, [78–79](#)

data cache, [78–79](#)

installing the DBMS, [78–79](#)

program cache, [79](#)

Memory, system performance

- caches, [328–330](#)
- consumption sources, [330–331](#)
- data cache, [329–330](#), [332–333](#)
- data integrity, [413–414](#)
- database log cache, [330](#)
- DB2 EDM pools, [335](#)
- estimating sufficient, [331–332](#)
- Internet structure cache, [330](#)
- procedure cache, [329–330](#), [335](#)
- sort cache, [330](#)
- MEMUSAGE option, [413–414](#)
- Merge-scan joins, [379–380](#)
- Message queuing software, [624–625](#)
- Messaging software, [624–625](#)
- Metadata
 - active sources, [690](#)
 - business metadata, [689](#)
 - data, definition, [686](#)
 - data dictionaries, [695–696](#)
 - data stewardship, [688](#)
 - data warehouse, [654](#)
 - data warehousing, [688](#)
 - definition, [16](#), [488](#)
 - examples, [16–17](#)
 - information, definition, [687](#)
 - integrated sources, [690](#)
 - knowledge, definition, [687](#)
 - nonsubvertible sources, [690](#)
 - overview, [685–686](#)
 - repositories, [691–695](#)
 - sources for, [690](#)
 - strategy for, [687–688](#)
 - system catalog, [689–691](#)

- technology metadata, [689](#)
- types of, [689–691](#)
- Microsoft Corporation
 - DBMS vendor, 63, [762](#), [767](#)
 - tool vendor, [772](#)
- Middleware, [192–193](#)
- Migration. *See* [Upgrading the DBMS](#).
- Mirror tables, [356](#)
- Mirroring, [597](#)
- MMDBMSs (main-memory database management systems), [596](#)
- Mobile DBMS architectures, [70](#)
- Mobile platforms, effect on DBAs, [53–55](#)
- ModelRight, [772](#)
- mongoDB, [56](#), [766](#)
- Monitoring
 - data warehouse performance, [652](#)
 - SP_MONITOR procedure, [345–346](#)
 - system-level performance, [345–346](#)
- Monitoring, database performance
 - contention, [23](#)
 - database administration tasks, [22–23](#)
 - factors affecting, [22–23](#)
 - optimization, [23](#)
 - resources, [22](#)
 - throughput, [22](#)
 - tools for, [313](#)
 - workload, [22](#)
- Monotonic page splits, [358–359](#)
- Moving data. *See also* [EXPORT utility](#); [IMPORT utility](#); [Loading data](#); [Unloading data](#).
 - backup, [535](#)
 - bulk movement, [623–625](#)
 - data warehouses, [644–645](#)

EXPORT utility, [622–623](#)
IMPORT utility, [622–623](#)
to multiple databases. See [Distributed databases](#).
tool vendors, [773](#)

MTBF (mean time between failures)
availability, [273–274](#)
disk drives, [580](#)

Mullins, Craig, Web site, [780](#)

Multi-index access, [387](#)

Multiple platforms, strategies for, [61–62](#)

Multitier implementation, [669–670](#)

MySQL, 64, [765](#)

N

Naming conventions

attributes, [116–119](#)

databases, [93–96](#)

devices, [364](#)

entities, [113](#)

NAS (network-attached storage), [605](#), [606](#)

Nested triggers, [343–344](#), [429](#)

Nested-loop joins, [379–380](#)

.NET framework, [194–195](#), [198](#)

NetIQ, [771](#)

Network data models, [754–755](#)

Network protocol architecture drivers, [673–674](#)

Network traffic

availability problems, [275](#)

connection pooling, [674](#)

database drivers, [672–674](#)

database gateways, [671–672](#)

performance problems, [670–671](#)

sniffing, [494–495](#)

- Newsgroups, [775–776](#)
- NHibernate, ORM library, [200](#)
- 99.999% availability, [273–274](#)
- NIST Web site, [782](#)
- Nodes, [155](#)
- Noncritical applications, [563](#)
- Noninvasive auditing methods, [494](#)
- Nonmatching index scan, [385](#)
- Nonstandard database objects, [94](#)
- Nonsubvertible metadata sources, [690](#)
- Normal forms
 - 1NF (first normal form), [129](#)
 - 2NF (second normal form), [129](#), [131](#), [132](#)
 - 3NF (third normal form), [132–133](#)
 - 4NF (fourth normal form), [134](#)
 - 5NF (fifth normal form), [134](#)
 - atomic, [129](#)
 - BCNF (Boyce Codd normal form), [134](#)
 - examples, [130–131](#)
- Normal page splits, [358–359](#)
- Normalization
 - definition, [128](#)
 - goals of, [128](#)
 - Internet databases, [679](#)
 - mapping logical to physical, [135](#)
- Normalized data model, [133–134](#)
- NoSQL
 - for availability, [296](#)
 - data models, [756](#)
 - DBMS vendors, [765–766](#)
 - effect on DBAs, [55–56](#)
- Nouns as
 - attributes, [124](#)

- entities, [115](#), [124](#)
- Nullability, specifying, [144](#)
- Nulling out, data masking, [498](#)
- Nulls
 - attributes, [119–120](#)
 - check constraints, [423–426](#)
 - loading data, [616](#)
- Number and date variance, data masking, [497](#)

O

- Obfuscation. See [Data masking and obfuscation](#).
- Object definitions, backup, [536–537](#)
- Object Management Group, [783](#)
- Object migration, tools for, [704–705](#)
- Object orientation
 - data models, [754–755](#)
 - DBMS vendors, [766](#)
 - ORM (object-relational mapping), [200](#)
 - relational databases, [199–200](#)
 - SQL, [199–200](#)
 - SQL (Structured Query Language), [199–200](#)
- Object Store (Progress Software), [766](#)
- Object-relational mapping (ORM), [200](#)
- Obsessive-Compulsive Data Quality blog, [781](#)
- ODBC (Object Database Connectivity)
 - callable routines, [192](#)
 - definition, [192](#)
 - drivers, [192](#), [673](#)
 - overview, [192–193](#)
- Off-loading logs, [339](#), [529](#)
- Offset tables, [588–590](#), [592](#)
- Off-site disaster recovery, [547](#)
- Off-site locations, disaster planning, [564](#)

OLAP (online analytical processing), [639](#), [640](#)
The OLAP Council, [783](#)
OLAP tool vendors, [773](#)
OLE DB, [193](#)
100 Year Archive Requirements Survey, [503n](#)
Online database reorganization, [288–289](#)
Online resources. *See* [Internet resources](#).
Ontos, [766](#)
Open database objects, system performance, [336](#)
The Open Group, [783](#)
Open-source software
 belief system, [764–765](#)
 choosing, [64](#)
 DBMS vendors, [764](#)
 definition, [764](#)
 vendors, [764–765](#)
Operating system failure, availability problems, [279](#)
Operating system support, strategies for, [65](#)
Operational support, standards, [102](#)
Operations, data models, [754](#)
Operations control, authorization, [467](#)
Optimizer, [374](#)
Optimizing database performance. *See also* [Database performance](#); [Relational optimization](#); [Reorganizing databases](#).
 cost-based *versus* rule-based, [344](#)
 database performance, managing, [305](#)
 performance monitoring and tuning, [23](#)
 split tables, [356](#)
Optimizing database performance, techniques for
 block size, [364–365](#)
 clustering, [356–358](#)
 combined tables, [356](#)
 compression, [361–362](#)

- database log placement, [363](#)
- denormalizing tables, [355–356](#)
- derivable data, [356](#)
- disk allocation, [364](#)
- distributed data placement, [363–364](#)
- file placement and allocation, [362–364](#)
- free space, [360–361](#)
- indexing, [352–355](#)
- interleaving data, [360](#). *See also* [Clustering](#).
- mapping physical files to database tables, [350–352](#)
- mirror tables, [356](#)
- monotonic page splits, [358–359](#)
- normal page splits, [358–359](#)
- overview, [349–350](#)
- page size, [364–365](#)
- page splitting, [358–359](#)
- parallelism, [351](#)
- partitioning, [350–352](#)
- physical denormalization, [356](#)
- prejoined tables, [355](#)
- raw partition *versus* file system, [351–352](#)
- redundant data, [356](#)
- repeating groups, [356](#)
- report tables, [355](#)
- speed tables, [356](#)
- split tables, [356](#)

OR logical operator, [403](#)

Oracle Corporation

- DBMS vendor, 63, [762](#)
- nonstandard database objects, [94](#)
- vendor contact, 63, [772](#)

Oracle Corporation (Hyperion), [773](#)

Oracle Magazine, [779](#)

Oracle program

blogs, [780](#)

IOUG (Independent Oracle Users Group), [740](#)

Web site, [778](#)

Oracle program, architecture

background processes, [326](#)

CKPT (checkpoint) process, [326](#)

config.ora file, [325](#)

control files, [325](#)

data files, [325](#)

database files, [325](#)

databases, [325](#)

DBWR (database writer) process, [326](#)

file categories, [325](#)

instances, [325](#)

LGWR (log writer) process, [326](#)

Oracle processes, [326](#)

overview, [325–327](#)

parameter files, [325](#)

PGA (program global area), [326](#)

physical structures, [325](#)

PMON (process monitor) process, [326](#)

processes, [326](#)

RECO (recover) process, [326](#)

redo log buffer, [326](#)

redo log files, [325](#)

server processes, [326](#)

SGA (system global area), [326](#)

SMON (system monitor) process, [326](#)

sort area, [326](#)

user processes, [326](#)

Oracle transportable tablespaces, [625](#)

ORDER BY clause, [388](#)

Ordered indexes, [157](#)
Ordering columns, [146](#)
Organization type, effect on DBMS strategy, [65](#)
Organizational design review, [237](#)
ORM (object-relational mapping), [200](#)
“Out of space” log conditions, [339–341](#)
Outer table, [379](#)
Overloading indexes, [355](#)

P

Page header
 corruption, [411](#)
 data page layouts, [588–589](#)
Page pointer, data page layouts, [592](#)
Page size, optimizing database performance, [364–365](#)
Page splits
 optimizing database performance, [358–359](#)
 reorganizing databases, [366](#)
Pages, recovering, [553](#)
Paradox, [767](#)
Parallel access, [390–391](#)
Parallelism
 CPU, [391](#)
 I/O, [390–391](#)
 optimizing database performance, [351](#)
 system, [391](#)
Parameter files, Oracle, [325](#)
Parent tables, [433–435](#)
Pareto (80/20) rule, [302](#)
Parity bits, [597](#)
Parsing database logs, [493–495](#)
Partial recovery, [542–543](#)
Partition scans, [381–382](#)

Partitioned index, [157](#). *See also* [b-tree index](#).

Partitioning, [350–352](#)

Passwords. *See also* [Security](#).

changing, [453](#)

creating, [454](#)

definition, [452](#)

disabling, [453](#)

embedding in code, [479](#)

guidelines for, [454](#)

limiting, [455](#)

locking, [453](#)

Patches

critical updates, [480–481](#)

for Oracle databases, [480–481](#)

security alerts, [480–481](#)

PCI (Payment Card Industry) DSS (Data Security Standard), [485](#), [491](#)

PCTFREE parameter, [360](#)

Performance. *See also* [Applications](#), [performance](#); [Database performance](#); [System performance](#).

benchmarks, [65–66](#)

DBA staffing requirements, [38](#)

gains from upgrading the DBMS, [85](#)

management, tools for, [708–714](#)

monitoring, tools for, [709–710](#)

RAID levels, [603](#)

Performance analysts, DBAs as, [36](#)

Performance problems

client/server computing, [670–674](#)

distributed databases, [632–633](#)

network traffic, [670–671](#)

Permissions. *See* [Privileges](#).

Personal computing, effect on DBAs, [53–55](#)

Personal DBMS architectures, [70](#)

Personnel, disaster planning, [569](#)
PGA (program global area), [326](#)
Phantoms, [218](#)
Physical data
 dependence, [376](#)
 independence, [757](#)
 modeling, [125–128](#)
Physical denormalization, [356](#)
Physical design review, [236](#)
PII (Personally Identifiable Information), [497](#)
PIT (point-in-time) recovery, [542–543](#), [545](#)
Planning
 change management, [245](#)
 for disaster. *See* [Disaster planning](#).
 outages, [286–287](#)
 SQL, [201](#)
 storage capacity, [608–609](#)
Plans (DB2), [94](#)
PLAN_TABLE, [395–398](#)
PMON (process monitor) process, [326](#)
Poet, [766](#)
Pointer consistency, [29](#)
Pointers for very large objects, [410–411](#)
Point-in-time recovery, [27](#)
Ponemon Institute, [450](#)
PostgreSQL, [765](#)
Post-implementation design review, [239](#)
Powersoft, [772](#)
Pre-implementation design review, [239](#)
Prejoined tables, [355](#)
Prepositional phrases, as attributes, [124](#)
Presentation logic, [664–666](#)
Primary key constraints, [416–417](#)

Primary key perspective, [435–436](#)

Primary keys

- database design, [144](#)

- indexes, [151](#)

Print servers, database connectivity, [664](#)

Privacy. *See* [Data, privacy](#).

Privacy Rights Clearinghouse, [449–450](#)

Privileged users, auditing, [495–496](#)

Privileges. *See also* [Authority](#); [Authorization](#); [Security](#).

- monitoring and reporting, [465](#)

- types of, [457–458](#)

Privileges, granting

- centralized administration, [457](#)

- database object privileges, [459](#)

- DCL (Data Control Language), [456–457](#)

- decentralized administration, [457](#)

- overview, [456–457](#)

- procedure privileges, [460](#)

- program privileges, [460](#)

- to PUBLIC authority, [460–461](#)

- system privileges, [459–460](#)

- table privileges, [458–459](#)

Privileges, revoking

- cascading REVOKEs, [462](#), [468](#)

- chronology and REVOKEs, [462–463](#)

- overview, [461](#)

Proactive performance, [306](#)

Proactivity, change management, [245](#)

Procedural DBAs

- duties of, [49](#)

- effect on DBAs, [46–50](#)

- managing database logic, [46](#)

- procedural database objects, [48](#), [49](#)

- role of, [49–50](#)
- stored procedures, [47](#), [48–50](#)
- triggers, [47](#), [48–50](#)
- UDFs (user-defined functions), [47](#), [48–50](#)

Procedure privileges, [460](#)

Procedures (programmatic). *See* [Triggers](#).

Procedures (standards). *See* [Standards and procedures](#).

Process monitor (PMON) process, [326](#)

Processes, Oracle, [326](#)

Professional advancement, DBA rule of thumb, [747–748](#)

Professional Association for SQL Server, [740](#), [783](#)

Professional certification. *See* [Certification](#).

Profiling data, [489](#)

Program global area (PGA), [326](#)

Program privileges, [460](#)

Programming and development, tools for, [724–726](#)

Progress Software, [766](#)

Propagation, [623–624](#), [722](#)

PUBLIC authority, [460–461](#)

Pure Java drivers, [673–674](#)

Purging data, data warehouse, [655](#)

Purging databases, *versus* archiving, [501](#)

The Pythian Group, [779](#), [781](#)

Q

Qualifying rows, [380](#)

Quality assurance testing, [101](#)

Quality of data. *See* [Data](#), [quality](#); [Data integrity](#).

Queries

- analysis, [378–379](#)
- performance, [650](#)
- tools for, [723–724](#)
- XML data, [203–205](#)

Query rewrite, [392–393](#)

Quest Software, [771](#)

Quiesce point, [528](#)

QUIESCE utility, [528](#)

R

RAC (Real Application Clusters), [294](#)

RAD (rapid application development), [677](#)

RAID (redundant array of independent disks)

definition, [597](#)

disk striping, [597](#)

error correction coding, [599](#)

fault tolerance, [601–602](#)

mirroring, [597](#)

parity bits, [597](#)

storage type, choosing, [603–604](#)

striping, [597](#)

RAID levels

performance, [603](#)

RAID-0, [597–598](#)

RAID-0+1, [602](#)

RAID-1, [598](#)

RAID-2, [599](#)

RAID-3, [599](#)

RAID-4, [600](#)

RAID-5, [600–601](#)

RAID-6, [601](#)

RAID-10, [601](#)

RAID-50, [602](#)

Raw files, database design, [149](#)

Raw partitions *versus* file systems, [351–352](#), [586–587](#)

RDA (Remote Database Access), [629–630](#)

Read efficiency, [333–335](#)

Read locks, [212–213](#)

RECO (recover) process, [326](#)

Recover to current, [26–27](#)

RECOVER utility, [553–554](#)

Recoverability

- availability, [267](#)
- goals of, [509–510](#)

Recover-to-current recovery, [541](#)

Recovery. *See also* [Backup](#); [Disaster planning](#).

- availability problems, [284](#)
- basic steps, [540–541](#)
- broken blocks or pages, [553](#)
- bulk-logged, [540](#)
- COBIT, [509–510](#)
- common reasons for, [548](#)
- data warehouse, [656–657](#)
- database administration tasks, [26–27](#)
- database logs, [338](#), [340](#)
- designing for, [533–534](#)
- dropped database objects, [552–553](#)
- duration of, [518](#), [549](#)
- full, [540](#)
- importance of, [515–516](#)
- indexes, [550–551](#)
- models, [340](#)
- objects, [534](#)
- optimum strategy, [547–549](#)
- options, determining, [538–539](#)
- overview, [537–538](#)
- planning for, [551](#)
- point-in-time recovery, [27](#)
- recover to current, [26–27](#)
- RECOVER utility, [553–554](#)

- regulatory compliance, [508](#)
- RMAN (Recovery Manager), [525–526](#)
- rows, [534](#)
- simple, [540](#)
- SQL Server models, [540](#)
- test databases, populating, [553–554](#)
- testing your plan, [551](#)
- tools for, [714–715](#)
- transaction recovery, [27](#)
- types of, [26–27](#)
- UNLOAD utility, [553–554](#)

Recovery, alternatives to

- disk mirroring, [556–557](#)
- redundant data, [555–556](#)
- replication, [555–556](#)
- snapshot replication, [555–556](#)
- standby databases, [554–555](#)
- symmetric replication, [555–556](#)

Recovery, types of

- matching to failure type, [549](#)
- off-site disaster, [547](#)
- partial, [542–543](#)
- PIT (point-in-time), [542–543](#), [545](#)
- recover to current, [541](#)
- to a recovery point, [543–544](#)
- REDO, [545–547](#)
- selecting, [548](#)
- transaction, [544–545](#)
- UNDO, [545–546](#)

Recovery Manager (RMAN), [525–526](#)

Red Gate Software, [771](#)

Redman, Thomas C., [489](#)

Redo log buffer, [326](#)

Redo log files, [325](#)

REDO recovery, [545–547](#)

Redundant array of independent disks (RAID). See [RAID \(redundant array of independent disks\)](#).

Redundant data

- backup/recovery alternative, [555–556](#)

- database design, [168–169](#)

- optimizing database performance, [356](#)

Reference customers, [68](#)

Referential constraints, [28](#), [146–147](#), [433](#)

Referential integrity (RI). See [RI \(referential integrity\)](#).

Regulatory compliance. See also [Data governance](#); *specific regulations*.

- best practices, [509](#)

- change management, [261–262](#)

- COBIT, [509–510](#)

- a collaborative approach to, [486–488](#)

- costs of compliance, [485](#)

- costs of non-compliance, [488](#)

- DBA tasks, [26](#), [487–488](#)

- importance to DBAs, [487–488](#)

- overview, [483–485](#)

- prosecution for non-compliance, [485](#)

- recoverability, [509–510](#)

REINDEX option, [413](#)

Relational Architects, [771](#)

Relational closure, [189–191](#)

Relational data models, [754–755](#)

Relational databases

- application design issues, [373–374](#)

- object orientation, [199–200](#)

Relational nulls, [423–426](#)

Relational optimization

- CPU costs, [376](#)

- database statistics, [376–378](#)
- definition, [375](#)
- density, [377](#)
- design issues, [374](#)
- duplicate values, [377](#)
- I/O costs, [376](#)
- joining tables, [379–381](#)
- optimizer, [374](#)
- physical data dependence, [376](#)
- query analysis, [378–379](#)
- query rewrite, [392–393](#)
- rule-based optimization, [393–394](#)
- view access, [391–392](#)
- view materialization, [392](#)
- view merging, [392](#)

Relational optimization, access path choices

- absolute positioning, [383–385](#)
- avoiding sorts, [387–388](#)
- CPU parallelism, [391](#)
- direct index lookup, [383](#)
- forcing, [398–399](#)
- hashed access, [389–390](#)
- index covering, [386–387](#)
- index screening, [386](#)
- indexed access, [382–389](#)
- index-only access, [386–387](#)
- I/O parallelism, [390–391](#)
- matching index scan, [383–385](#)
- multi-index access, [387](#)
- nonmatching index scan, [385](#)
- parallel access, [390–391](#)
- partition scans, [381–382](#)
- relative positioning, [385](#)

- reviewing, [394–398](#)
- system parallelism, [391](#)
- table scans, [381–382](#)
- tablespace scans, [381–382](#)

Relations, data models, [755](#)

Relationships, [122–123](#)

Relationships between entities

- cardinality, [122–123](#)
- definition, [122](#)
- degree. *See* [Cardinality](#).
- description, [122](#)
- discovering, [124–125](#)
- optionality, [123](#)
- verbs as, [124](#)

Relative positioning, [385](#)

Release schedules, effect on DBMS strategy, [68](#)

Release upgrades, backup, [534](#)

Releases *versus* versions, [82–87](#)

Reliability, availability, [267](#)

Remote Database Access (RDA), [629–630](#)

Remote mirroring, disaster planning, [573](#)

Remote requests, [630–631](#)

Remote unit of work, [630–631](#)

Removing. *See* [Deleting](#).

REORG, [288–289](#)

REORG utility, [368–369](#)

Reorganizing databases. *See also* [Database performance](#); [Optimizing database performance](#).

- automation, [371](#)
- causes of disorganization, [365–369](#)
- cluster ratios, [369](#)
- determining the need for, [369](#)
- disorganized tablespace, [367–368](#)

- ensuring availability, [288–289](#)
- file extents, [366](#)
- fragmentation, [366](#)
- gathering statistics, [370](#)
- indexes, [369–370](#)
- leaf distance, [370](#)
- manually, [368–369](#)
- online, [288–289](#)
- page splits, [366](#)
- row chaining, [366](#)
- row migration, [366](#)
- tools for, [314](#)
- unclustered data, [366](#)
- utilities for, [368–369](#)

REPAIR utility (DB2), [411–412](#)

REPEATABLE READ isolation, [217](#)

Repeating groups

- denormalization, [169–170](#)
- optimizing database performance, [356](#)

Replication

- backup/recovery alternative, [555–556](#)
- bulk data movement, [623–624](#)
- tools, [722](#)

Report tables, [355](#)

Reporting, tools for, [723–724](#)

Repositories, [691–695](#)

Repository Manager (IBM), [695](#)

Repository tools, vendors, [772–773](#)

Required applications, [563](#)

Resources

- effective use of, DBA rule of thumb, [745–746](#)
- performance factor, [301](#)
- performance monitoring and tuning, [22](#)

Response time, [266](#)

Responsive Systems, [771](#)

REST (representational state transfer), [681](#)

Reverse key index, [156–157](#)

Reviewing access path choices, [394–398](#)

REVOKE statements, [456–457](#), [461](#)

RI (referential integrity)

- versus* check constraints, [441–442](#)
- child tables, [433–434](#)
- DBMS support for, [438](#)
- definition, [433](#)
- deleting rows, [435–436](#)
- foreign key perspective, [434–435](#)
- foreign key values, [434–436](#)
- overview, [146–147](#), [433–434](#)
- parent tables, [433–435](#)
- primary key perspective, [435–436](#)
- versus* program logic, [441–442](#)
- referential constraints, [433](#)
- relationships, setting up, [436–437](#)
- rules, [434–436](#)
- rules of thumb, [442–444](#)
- self-referencing constraints, [437](#)
- system-managed, [441](#)
- tools for, [705](#)
- with triggers, [438–441](#)
- user-managed, [441](#)

Riak, [56](#)

Richard Foote's Oracle Blog, [780](#)

Rigas, John, [485](#)

Rigas, Tony, [485](#)

Rightsizing, and database connectivity, [662](#)

Risk, disaster planning

- assessing, [561–563](#)
- business service interruption, [561–563](#)
- categories of, [561](#)
- financial loss, [561–563](#)
- legal responsibilities, [561–563](#)

Risk management

- tools for, [716–721](#)
- upgrading the DBMS, [84–86](#)

RMAN (Recovery Manager), [525–526](#)

Rocket Software, [771](#)

Roles and responsibilities

- authorization by, [466](#), [468](#)
- standards, [96–98](#), [97](#)

Ross E/R method, [112](#)

Row-level triggers, [432](#)

Rows

- chaining, [366](#)
- data page layouts, [590](#)
- deleting, [435–436](#)
- headers, data page layouts, [590](#)
- length, data page layouts, [592](#)
- migration, [366](#)
- recovering, [534](#)
- size, specifying, [148](#)

Ruby on Rails, [198](#)

Rule-based optimization

- versus* cost-based, [344](#)
- relational optimization, [393–394](#)

Rules. *See also specific rules.*

- check constraints, [424](#)
- definition, [94](#)
- referential integrity, [434–436](#)
- standards, [94](#)

Rumbaugh, James, [113](#)
Rumbaugh E/R method, [112](#)
RUNSTATS utility, [377–378](#)

S

SA (system administration), [20](#), [21](#)
Salaries of DBAs, [4–6](#)
SAN (storage area network), [278](#), [604–605](#), [606](#)
SAP (Business Objects), [773](#)
Sarbanes-Oxley (SOX) Act, [483](#), [485](#), [491](#)
SAS Institute, [773](#)
Scalability
 data warehouse, [649](#)
 effect on DBMS strategy, [66](#)
Scribe, design review, [230](#)
Scripts, change management, [262](#)
SCSI (small computer system interface), [605](#)
SearchDataManagement portal, [781](#)
SearchOracle portal, [781](#)
SearchSQLServer portal, [781](#)
Secured Hash Algorithm (SHA-1), [472](#)
Security. *See also* [Authority](#); [Encryption](#); [Passwords](#); [Privileges](#).
 administrator authorization, [467](#), [468](#)
 auditing, [477–478](#)
 authentication, [452](#)
 availability problems, [280](#)
 basics, [451–455](#)
 centralizing, [26](#)
 costs of data breaches, [450](#)
 data breaches, [449–450](#)
 data theft (example), [496](#)
 database administration tasks, [24–25](#)
 database users, [455–456](#)

- external threats, [478–480](#)
- fixpacks and maintenance, [480–481](#)
- horizontal restriction, [469](#)
- job scheduling, [479](#)
- logic-oriented, [470](#)
- non-DBA, [480](#)
- options for system performance, [344](#)
- replacement tools for, [721](#)
- scope of the problem, [449–450](#)
- sensitive data sets, [478–479](#)
- standards, [100–101](#)
- with stored procedures, [470](#)
- tools for, [720–721](#)
- user names, [456](#)
- using views for, [468–470](#)
- vertical restriction, [469](#)

Security, login

- administration, rules of thumb, [453](#), [455](#)
- definition, [455–456](#)
- limiting, [455](#)
- required information, [452–453](#)

Security, passwords

- changing, [453](#)
- creating, [454](#)
- definition, [452](#)
- disabling, [453](#)
- embedding in code, [479](#)
- guidelines for, [454](#)
- limiting, [455](#)
- locking, [453](#)

SEGUS Inc., [771](#)

SEI Web site, [782](#)

SELECT INTO statements, [341](#)

SELECT INTO/BULKCOPY option, [341](#)
SELECT privileges, [458–459](#)
Selective auditing methods, [494](#)
Self-referencing constraints, [437](#)
Semantic data integrity. *See* [Data integrity](#), [semantic](#).
Semantic integrity, [28](#)
Serena, [773](#)
SERIALIZABLE isolation, [218](#)
Server processes, [326](#)
Servers. *See also* [Client/server computing](#); [SQL Server \(Microsoft\)](#).
 application, [664](#)
 availability problems, [276](#), [283–284](#)
 database, [664](#)
 definition, [665–666](#)
 file, [664](#)
 hardware failure, availability problems, [276](#)
 list, [777](#)
 performance, [283–284](#), [650](#)
 print, [664](#)
 transaction, [207–210](#)
Servers, database
 definition, [664](#)
 hosting, [675](#)
 location, upgrading, [88](#)
Service level agreements (SLAs), [38](#)
Serviceability, availability, [268](#)
Service-level management (SLM), [308–311](#)
Service-oriented architecture (SOA), [680](#)
Servlets, [196–197](#)
Set theory, SQL, [190](#)
Set-at-a-time processing, [189–191](#)
SGA (system global area), [326](#)
SHA-1 (Secured Hash Algorithm), [472](#)

SHA-256 hashing, [472](#)
Shared locks, [212–213](#)
Shared-disk clustering, [72](#), [294](#)
Shared-nothing clustering, [71–72](#), [294](#)
Sharing knowledge, DBA rule of thumb, [739–741](#)
SHOWPLAN command, [394–398](#), [712](#)
Shuffling, data masking, [497](#)
Silos, in a fractured environment, [310–311](#)
Simple Object Access Protocol (SOAP), [680](#)
SIMPLE parameter, [620–621](#)
Simple recovery, [340](#), [540](#)
Simplification, DBA rule of thumb, [741–742](#)
Size terminology, data storage, [582](#)
Skilling, Jeff, [485](#)
SKIP LOCKED DATA parameter, [219](#)
Skipping
 interim releases, [86–87](#)
 locked rows, [219](#)
SLAs (service level agreements), [38](#)
SLM (service-level management), [308–311](#)
Small computer system interface (SCSI), [605](#)
SMON (system monitor) process, [326](#)
Snapshot replication, [555–556](#)
SNIA (Storage Networking Industry Association), [503n](#), [783](#)
Snowflake schema, data warehouse, [643](#)
SOA (service-oriented architecture), [680](#)
SOAP (Simple Object Access Protocol), [680](#)
Softbase Systems Inc., [771](#)
Software AG, [764](#)
Software environment, designing, [193–194](#)
SoftwareOnZ LLC, [771](#)
Solid state devices (SSDs), [323–324](#), [596](#)
SolidDB (IBM), [596](#)

Sort area, [326](#)

Sorting

avoiding, [387–388](#)

indexes, [152](#)

SQL tuning, [404](#)

tools for, [314](#)

SOX (Sarbanes-Oxley) Act, [483](#), [485](#), [491](#)

Space management. *See also* [Storage management](#).

monitoring usage, [587–588](#)

tools for, [726](#)

Space management, data page layouts

allocation pages, [589](#)

allocation units, [589](#)

bitmaps, [589](#)

data record layouts, [590](#)

data rows, [588–589](#)

header information, [592](#)

index key values, [592](#)

index page layouts, [592–594](#)

offset and adjust tables, [592](#)

offset table, [588–589](#)

offset tables, [590](#)

overview, [588–589](#)

page header, [588–589](#)

page pointer, [592](#)

row data, [590](#)

row header, [590](#)

row length, [592](#)

sample, [589](#)

space page map, [589](#)

table size, calculating, [591–592](#)

transaction logs, [594–595](#)

Space page map, data page layouts, [589](#)

- SP_CONFIGURE procedure, [80n.4](#)
- Speed tables, optimizing database performance, [356](#)
- Split tables
 - horizontally split, [166](#)
 - overview, [165–166](#)
 - vertically split, [166](#)
- Splitting text columns, [166–168](#)
- SP_MONITOR procedure, [345–346](#)
- SPUFI, [81n.5](#)
- SQData, [773](#)
- SQL (Structured Query Language)
 - access paths, [187](#)
 - APIs, [192–193](#)
 - benefits of, [188](#)
 - binding, [477](#)
 - callable routines, [192](#)
 - code, design review, [238](#)
 - coding for performance, [202–203](#)
 - COM, [193](#)
 - creating with code generators, [191–192](#)
 - cursor, [190](#)
 - definition, [186–187](#)
 - dynamic, [201](#)
 - embedded, [201](#)
 - embedding in programs, [191–192](#)
 - JDBC (Java Database Connectivity), [192–193](#)
 - joins, [189](#)
 - middleware, [192–193](#)
 - object orientation, [199–200](#)
 - ODBC (Object Database Connectivity), [192–193](#)
 - OLE DB, [193](#)
 - overview, [186–188](#)
 - planned, [201](#)

- query analysis, [378–379](#), [713](#)
- query rewrite, [392–393](#)
- querying XML data, [203–205](#)
- relational closure, [189–191](#)
- set theory, [190](#)
- set-at-a-time processing, [189–191](#)
- SQL/XML, [204](#)
- standalone, [201](#)
- standards Web site, [782–783](#)
- static, [201](#)
- subqueries, [189](#)
- syntax, [187](#)
- types of, [200–201](#)
- unplanned, [201](#)
- usage considerations, *188*, *202*
- XQuery language, [204](#)

SQL injection attacks

- examples, [474](#)
- overview, [201–202](#), [473–475](#)
- preventing, [475–476](#)
- static *versus* dynamic SQL, [476](#)

SQL Is Your Friend blog, [780](#)

SQL Marklar blog, [780](#)

SQL Rockstar blog, [780](#)

SQL Server (Microsoft)

- filegroups, [149](#)
- nonstandard database objects, [94](#)
- Professional Association for SQL Server, [740](#)
- transaction logs, backup, [530](#)
- vendor contact, *63*
- Web site, [778](#)

SQL Server Pro, [779](#)

SQL tuning

- basic steps, [399–400](#)
- Cartesian products, [402](#)
- code generators, [405](#)
- COMMIT frequency, [404–405](#)
- finding problem statements, [303–304](#), [406–407](#)
- LIKE logical operator, [403–404](#)
- OR logical operator, [403](#)
- overview, [202–203](#)
- rules of thumb, [400–406](#)
- sorts, [404](#)
- stored procedures, [405–406](#)
- tools for, [313](#)

SQL...User Group portal, [782](#)

SQL/XML, [204](#)

SSDs (solid state devices), [323–324](#), [596](#)

Standalone SQL, [201](#)

Standardizing default values, data warehouse, [647](#)

Standards and procedures

- abbreviations, [96](#)
- application development, [100](#)
- clusters, definition, [94](#)
- communications, [98](#)
- data administration, [98–99](#)
- database links, definition, [94](#)
- database naming conventions, [93–96](#)
- definition, [93](#)
- design review guidelines, [102](#)
- distributed data, [629–630](#)
- filegroups, definition, [94](#)
- importance of, [93](#)
- Internet resources, [782](#)
- migration and turnover, [101–102](#)
- nonstandard database objects, [94](#)

- online manuals, [727–728](#)
- operational support, [102](#)
- plans, definition, [94](#)
- roles and responsibilities, [96–98](#), 97
- rules, definition, [94](#)
- security, [100–101](#)
- storage groups, definition, [94](#)
- system administration, [100](#)

Standby databases

- availability problems, [276](#), [277](#)
- backup/recovery alternative, [554–555](#)
- versus* backups, [277](#)
- DB2 HADR (high-availability disaster recovery), [285](#)
- definition, [276](#)
- disaster planning, backup, [573](#)
- Oracle, [277](#)

Star schema, data warehouse, [641–643](#)

Statement-level triggers, [432](#)

Static SQL, [201](#)

Storage. *See specific media.*

Storage area network (SAN), [278](#), [604–605](#), [606](#)

Storage groups, [94](#)

Storage management. *See also* [Files and data sets](#); [Space management](#); *specific media.*

- capacity planning, [608–609](#)
- cool data, [607–608](#)
- dormant data, [607–608](#)
- fragmentation, [595](#)
- goals for, [583](#)
- hot data, [607–608](#)
- integrity *versus* availability, [580](#)
- multitemperature data, [607–608](#)
- overview, [579–583](#)

rate of data growth, [581–582](#)

size terminology, [582](#)

warm data, [607–608](#)

Storage management, media options. *See also specific media.*

disk, [596](#)

fiber channel, [605](#)

JBOD (just a bunch of disks), [604](#)

MMDBMSs (main-memory database management systems), [596](#)

NAS (network-attached storage), [605](#), [606](#)

overview, [596](#)

SAN (storage area network), [604–605](#), [606](#)

SCSI disks, [605](#)

SSDs (solid state devices), [596](#)

tape, [596](#)

tiered storage, [606–608](#)

Storage management, software

backup/recovery alternative, [535–536](#), [547](#)

disaster planning, backup, [572–573](#)

Storage Networking Industry Association (SNIA), [503n](#), [783](#)

Storage requirements

database design, [148](#)

installing the DBMS, [76–78](#)

Stored procedures

procedural DBAs, [47](#), [48–50](#)

as security tools, [470](#)

SQL tuning, [405–406](#)

Strategies for DBMS

benchmarks, TPC, [65–66](#)

choosing a DBMS, [63–68](#). *See also* [Vendors](#), [DBMS](#).

cloud database systems, [74](#)

cost of ownership, [67](#)

DBMS architectures, [68–71](#)

DBMS clustering, [71–73](#)

- DBMS proliferation, [73](#)
- factors affecting, [65–68](#)
- hardware issues, [73–74](#)
- multiple platforms, [61–62](#)
- operating system support, [65](#)
- organization type, [65](#)
- product complexity, [68](#)
- reference customers, [68](#)
- release schedules, [68](#)
- scalability, [66](#)
- technical support, [67](#)
- tool availability, [66](#)

Striping, [597](#)

Structural data integrity. *See* [Data integrity, database structure](#).

Structure, data models, [754](#)

Structured Query Language (SQL). *See* [SQL \(Structured Query Language\)](#).

Subqueries, SQL, [189](#)

Substitution, data masking, [497](#)

Subsystem failure, backup, [533](#), [550](#)

Suppliers. *See* [Vendors](#).

Support policies for old releases, [89](#)

Supporting *versus* exploiting, [91](#)

Swartz, Mark, [485](#)

Sybase Inc.

- ISUG (International Sybase User Group), [740](#), [783](#)

- vendor contact, 64, [763](#)

- Web site, [778](#)

Symmetric replication, [555–556](#)

Synonyms, in data modeling, [118](#)

Sysplex (IBM), [294–295](#)

System administration (SA), [20](#), 21

System administration standards, [100](#)

System administrators

- authorization, [467](#)
- limiting number of, [468](#)
- System catalog, [342–343](#), [689–691](#)
- System catalog tables, [345](#)
- System DBAs, [31–32](#)
- System global area (SGA), [326](#)
- System memory failure, [276](#)
- System monitor (SMON) process, [326](#)
- System monitoring, [345–346](#)
- System parallelism, [391](#)
- System performance
 - allied agents, [321–322](#)
 - DBMS components, [324](#)
 - disk storage and I/O, [322–324](#)
 - hardware configuration, [322–324](#)
 - operating system interaction, [320–321](#)
 - overview, [319–320](#)
 - tools for, [709–710](#)
- System performance, DBMS installation and configuration
 - cache, [328–330](#)
 - configuration types, [327–328](#)
 - contention, [341–342](#)
 - data cache, [329–330](#), [332–335](#)
 - database log cache, [330](#)
 - database logs, [336–341](#)
 - deadlock detection, [341](#)
 - deadlocks, [342](#)
 - defaults, [344](#)
 - distributed database, [344](#)
 - guidelines, [344](#)
 - identity values, [344](#)
 - Internet structure cache, [330](#)
 - lock suspensions, [341](#)

- locking, [341–342](#)
- memory, [328–332](#)
- nested trigger calls, [343–344](#)
- open database objects, [336](#)
- optimization, cost-based *versus* rule-based, [344](#)
- procedure cache, [329–330](#), [335](#)
- read efficiency, [333–335](#)
- sample options, [343–344](#)
- security options, [344](#)
- sort cache, [330](#)
- system catalog, [342–343](#)
- time-outs, [342](#)

System privileges, [459–460](#)

System time, [179–180](#)

System-managed referential integrity, [441](#)

T

Table editors, [707–708](#)

Table scans, [151](#), [381–382](#)

Tables

- adjust, [592](#)
- combined, [168](#)
- dropping, [250–252](#)
- mirror, [165](#)
- naming conventions, [95–96](#)
- offset, [592](#)
- prejoined, [164](#)
- privileges, [458–459](#)
- report, [164–165](#)
- size, calculating, [591–592](#)
- size control, [585](#)
- speed, [172–173](#)
- split, [165–166](#)

- storage requirements, calculating, [590–592](#)
- Sybase segments, [585](#)
- Tablespace
 - database design, [148](#)
 - disorganized, [367–368](#)
 - scans, [381–382](#)
- Table-to-table synchronization, [498](#)
- Tamino (Software AG), [764](#)
- Tape storage. *See also* [Storage management](#).
 - DBMS requirements, [77](#)
 - disaster planning, backup, [570–571](#)
 - storage management option, [596](#)
 - WORM (write once, read many) technology, [596](#)
- Tapping requests, [494–495](#)
- Task-oriented DBAs, [36](#)
- TBCHECK utility (Informix), [411](#)
- TDES (Triple DES), [472](#)
- Team members, disaster planning and recovery, [569](#)
- Technical education, DBA rule of thumb, [746–747](#)
- Technical support, effect on DBMS strategy, [67](#)
- Technology, effects on DBAs
 - Big Data movement, [55–56](#)
 - cloud computing, [53–55](#)
 - database-coupled application logic, [46–50](#)
 - Internet and e-business, [50–52](#)
 - managing database logic, [46–50](#)
 - mobile platforms, [53–55](#)
 - NoSQL, [55–56](#)
 - personal computing, [53–55](#)
 - procedural DBAs, [46–50](#)
- Technology metadata, [689](#)
- Technology silos, in a fractured environment, [310–311](#)
- TechTarget, [781](#)

Temporal data support, [177–180](#)
Temporal database systems, data integrity, [444–446](#)
Temporal requirements, database design, [177–180](#)
Teradata Corporation, 64, [607](#), [763](#)
Teradata Magazine, [779](#)
Test beds, loading/unloading data for, [621–622](#)
Test databases, populating, [553–554](#)
Testing
 disaster planning, recovery, [567–569](#), [574](#)
 recovery plans, [551](#)
 tools for, [725](#)
Thin clients, [670](#)
Throughput
 performance factor, [301](#)
 performance monitoring and tuning, [22](#)
Tibbetts, Hollis, [488](#)
Tier-1 DBMS vendors, 63, [762](#)
Tier-2 DBMS vendors, 64, [763](#)
Tiered storage, [606–608](#)
Time outs, locks, [213–214](#)
Time zones, availability across, [270–271](#)
Time-outs, system performance, [342](#)
TimesTen (Oracle), [596](#)
Tool vendors. *See* [Vendors](#), [tools](#).
Tools, for. *See also specific tools*.
 capacity planning, [313](#)
 catalog query and analysis, [705–707](#)
 checkpoint/restart, [725](#)
 compliance, [716–721](#)
 debugging, [726](#)
 end-to-end performance, [713–714](#)
 replication, [722](#)
 repositories, [772–773](#)

trending, [719](#)

utility management, [716](#)

Tools, for availability. *See also* [Standby databases](#).

AlwaysOn features, [285](#)

Database Definition on Demand, [289–290](#)

DB2 HADR (high-availability disaster recovery), [285](#)

effect on DBMS strategy, [66](#)

RAC (Real Application Clusters), [294](#)

REORG, [288–289](#)

Tools, for DBAs

analytics, [721–724](#)

application performance, [711–713](#)

auditing tools, [717–719](#)

availability, effect on DBMS strategy, [66](#)

backup and recovery, [714–715](#)

benefits of, [699–700](#)

business intelligence, [721–724](#)

catalog query and analysis, [705–707](#)

catalog visibility, [706](#)

change management, [254](#), [701–703](#)

checkpoint/restart, [725](#)

compliance, [716–721](#)

compression, [726–727](#)

cost justification, [702](#), [731](#)

data integrity, [411–414](#)

data masking, [720](#)

data modeling and design, [700–701](#)

data profiling, [719–720](#)

data protection, [716–721](#)

data warehousing, [721–724](#)

database comparison, [703–704](#)

database performance, [711](#)

database utilities, [715–716](#)

- DBA staffing requirements, [39](#)
- debugging, [726](#)
- end-to-end performance, [713–714](#)
- ETL (extract, transfer, load), [721–723](#)
- governance, [716–721](#)
- homegrown, [732–733](#)
- invasive performance tools, [710](#)
- native *versus* third-party, [728](#)
- object migration, [704–705](#)
- online standards manuals, [727–728](#)
- performance management, [708–714](#)
- performance monitor, [709–710](#)
- programming and development, [724–726](#)
- propagation, [722](#)
- query, [723–724](#)
- reorganizing databases, [368–369](#)
- replication, [722](#)
- reporting, [723–724](#)
- RI (referential integrity), [705](#)
- risk management, [716–721](#)
- security replacement, [721](#)
- security tools, [720–721](#)
- space management, [726](#)
- system performance, [709–710](#)
- table editors, [707–708](#)
- testing, [725](#)
- trending, [719](#)
- types of, [699–700](#)
- utility management, [716](#)
- vendors, evaluating, [729–732](#)

Tools, vendors for. *See also specific vendors.*

- business intelligence tools, [773](#)
- data modeling tools, [771–772](#)

- data movement tools, [773](#)
- DBA tools, [729–732](#), [770–771](#)
- ETL tools, [773](#)
- OLAP tools, [773](#)
- repository tools, [772–773](#)
- TP (transaction processing) system, [207–209](#)
- TPC (Transaction Processing Performance Council), [66–67](#)
- Trace-based auditing, [493–495](#)
- Tracker tables, [308](#)
- Transaction failure, backup, [516–517](#), [550](#)
- Transaction logs. *See also* [Database logs](#).
 - backing up, [530](#)
 - data page layouts, [594–595](#)
 - file placement, [585](#)
- The Transaction Processing Council, [783](#)
- Transaction processing monitor, [207–209](#)
- Transaction recovery, [27](#), [544–545](#)
- Transaction servers, [207–209](#)
- Transaction time, [179–180](#)
- Transactions
 - ACID properties, [205–206](#)
 - application servers, [209–210](#)
 - atomicity, [205–206](#)
 - consistency, [206](#)
 - definition, [205](#)
 - durability, [206](#)
 - example, [206–207](#)
 - guidelines, [207](#)
 - isolation, [206](#)
 - TP (transaction processing) system, [207–209](#)
 - transaction processing monitor, [207–209](#)
 - transaction servers, [207–209](#)
 - UOW (unit of work), [207](#)

Transition tables, [430–431](#)
Transition variables, [430–431](#)
Transparency, distributed databases, [626](#)
Transparent encryption, [473](#)
Transportable tablespaces (Oracle), [625](#)
Treehouse Software, [773](#)
Trending, tools for, [719](#)
Triggers
 active databases, [426](#)
 definition, [426](#)
 DELETE, [438–441](#)
 example, [431](#)
 firing, [428–429](#)
 firing while loading data, [615](#)
 granularity, [431–432](#)
 implementing referential integrity, [429–430](#)
 INSERT, [438–441](#)
 INSTEAD OF, [432](#)
 multiple on same table, [428–429](#)
 nested, [429](#)
 overview, [426–428](#)
 procedural DBAs, [47](#), [48–50](#)
 referential integrity, [438–441](#)
 row level, [432](#)
 semantic data integrity, [426–433](#)
 statement level, [432](#)
 transition tables, [430–431](#)
 transition variables, [430–431](#)
 UPDATE, [438–441](#)
Triple DES (TDES), [472](#)
TRUNC LOG ON CHKPT option, [340](#), [530](#)
TRUNCATE TABLE statements, [341](#)
Tuning performance. See [Database performance, tuning](#).

Turnover, standards, [101–102](#)

TUSC, [779](#)

24/24 availability, [270–271](#)

Twitter, as a resource, [741](#)

Two-phase COMMIT, [631](#)

Tyco, [485](#)

Type 1 drivers, [673](#)

Type 2 drivers, [673](#)

Type 3 drivers, [673–674](#)

Type 4 drivers, [673](#)

U

UDFs (user-defined functions), [47](#), [48–50](#)

UDT (user-defined data types), [418–419](#)

UML (Unified Modeling Language), [113](#), [114](#)

Unclustered data, reorganizing databases, [366](#)

UNCOMMITTED READ isolation, [216–217](#)

UNDO recovery, [545–546](#)

Unfederated multidatabase schemes, [627](#)

UNION clause, [388](#)

Unique constraints, [28](#), [417](#)

Unique entity identification, [416–417](#)

Unique identifiers for columns, [145](#)

Unit testing, database environment, [101](#)

UNLOAD utility, [553–554](#), [618–621](#), [621–622](#). *See also* [EXPORT utility](#).

Unloading data. *See also* [Loading data](#).

- for application test beds, [621–622](#)

- backup/recovery alternative, [534–535](#)

- concurrency, [619](#)

- encoding scheme, specifying, [620](#)

- floating-point data, [620](#)

- generating LOAD parameters, [620](#)

- from image copy backups, [619](#)

- LIMIT parameter, [620–621](#)
- limiting, [620–621](#)
- number of rows, specifying, [620–621](#)
- partial unload, [620–621](#)
- selection criteria, specifying, [620–621](#)
- SIMPLE parameter, [620–621](#)
- UNLOAD utility, [618–621](#), [621–622](#)
 - from views, [621](#)
 - WHEN clause, [620–621](#)
- Unplanned outages, [286–287](#)
- Unplanned SQL, [201](#)
- UOW (unit of work), [207](#)
- Update locks, [213](#)
- UPDATE privileges, [458–459](#)
- UPDATE rule, [434–436](#)
- UPDATE statements
 - modifying temporal data, [180](#)
 - in triggers, [429](#)
- UPDATE STATISTICS command, [377–378](#)
- UPDATE trigger, [438–441](#)
- Upgrading the DBMS. *See also* [Configuring the DBMS](#); [Installing the DBMS](#).
 - application complexity, [88](#)
 - benefits of, [83–84](#)
 - costs, [84](#), [85](#)
 - database server location, [88](#)
 - DBA staff skill set, [90](#)
 - DBMS environment complexity, [87–88](#)
 - deprecated features, [85n.7](#)
 - fallback planning, [92](#)
 - features and complexity, [87](#)
 - migration standards, [101–102](#)
 - migration verification, [92](#)

- organizational style, [89–90](#)
- overview, [82–87](#)
- performance gains, [85](#)
- platform support, [90–91](#)
- risks, [84–86](#)
- skipping interim releases, [86–87](#)
- strategy for, [92](#)
- support policies for old releases, [89](#)
- supporting software, [91](#)
- supporting *versus* exploiting, [91](#)
- vendor reputation, [89](#)
- versions *versus* releases, [82–87](#)

Upsizing, and database connectivity, [662](#)

U.S. Public...Protection Act of 2002, [483](#)

Usenet newsgroups, [775–776](#)

User acceptance testing, database environment, [101](#)

User processes, [326](#)

User-defined data types (UDT), [418–419](#)

User-defined functions (UDFs), [47](#), [48–50](#)

User-managed referential integrity, [441](#)

Users

- DBA staffing requirements, [38](#), [39](#)
- group associations, [740](#)
- names, [456](#)
- privileged, auditing, [495–496](#)
- security, [455–456](#)

Utility management, tools for, [716](#)

Utility programs. *See* [Tools](#); *specific programs*.

V

V\$ tables, [370](#)

Valid time, [179–180](#)

Variable-length columns, [144](#)

Vendors

- evaluating, [729–732](#)
- reputation, importance of, [89](#)

Vendors, DBMS

- Action Corporation, [764](#)
- Big Three, [762](#)
- dBase, [767](#)
- FileMaker, [767](#)
- IBM Corporation, [762](#)
- Informix, [763](#)
- Ingres, [763–764](#)
- Lotus, [767](#)
- main groups, [761–762](#)
- Microsoft, [762](#), [767](#)
- NoSQL systems, [765–766](#)
- object-oriented systems, [766](#)
- Ontos, [766](#)
- open-source systems, [764](#)
- Oracle Corporation, [762](#)
- PC-based systems, [766–767](#)
- Poet, [766](#)
- Progress Software, [766](#)
- Software AG, [764](#)
- Sybase Inc., [763](#)
- Teradata Corporation, [763](#)
- tier 1, 63, [762](#)
- tier 2, 64, [763](#)
- Web sites for, [778](#)

Vendors, tools. *See also specific vendors.*

- business intelligence tools, [773](#)
- data modeling tools, [771–772](#)
- data movement tools, [773](#)
- DBA tools, [729–732](#), [770–771](#)

- ETL tools, [773](#)
- OLAP tools, [773](#)
- repository tools, [772–773](#)
- Verifying the DBMS install, [81](#)
- Versions *versus* releases, [82–87](#)
- Vertical restriction, [469](#)
- Very critical applications, [562](#)
- View access, relational optimization, [391–392](#)
- View materialization, relational optimization, [392](#)
- View merging, relational optimization, [392](#)
- Views
 - description, [175–176](#)
 - as security tools, [468–470](#)
 - unloading data from, [621](#)
 - uses for, [176–177](#)
- Visible Systems, [772](#)
- Vision Solutions, [773](#)
- Visual Insights, [773](#)
- Volatility, DBA staffing requirements, [39](#)
- VPD (Virtual Private Database) (Oracle), [471](#)

W

- Wallets, [473](#)
- WAN (wide-area network), disaster planning, [573](#)
- Web Farming site, [783](#)
- Web resources. *See* [Internet resources](#).
- Web services. *See also* [Database connectivity](#); [Internet](#).
 - definition, [680](#)
 - REST (representational state transfer), [681](#)
 - SOA (service-oriented architecture), [680](#)
 - SOAP (Simple Object Access Protocol), [680](#)
- WHEN clause, [620–621](#)
- White Sands Technology, Inc., [771](#)

Wide-area network (WAN), disaster planning, [573](#)

Winter Corporation, [581–582](#)

WITH GRANT OPTION, [457](#)

Workload

- performance factor, [301](#)

- performance monitoring and tuning, [22](#)

WorldCom, [485](#)

Write locks, [212–213](#)

Write-ahead logs, [337](#)

X

XML (eXtensible Markup Language), [204](#)

XML data, querying, [203–205](#)

The XML portal, [783](#)

XQuery language, [204](#)

Y

Yevich, Lawson & Associates, [779](#)

Z

z/OS, data sharing, [295](#)

Footnotes

Chapter 1

1. Social media have made this trend even more pronounced. People are living their lives out loud and in a recorded, and therefore data-driven, format.
2. GridApp Systems study titled “Trends in Database Management: The Automation Opportunity,” 2006.
3. Dice.com is a career Web site for technology and engineering professionals and the companies that seek to employ them.
4. Global Knowledge is a technology training firm.
5. Tech Republic is a Web portal published by ZDNet.
6. Janco Associates, Inc., is a management consulting firm that focuses on Management Information Systems.
7. www.computerworld.com/s/article/9224243/IT_Salary_Survey_2012.
8. www.bls.gov/.
9. I refer to Adabas and Ingres as “fringe” not due to any functional or technical deficiencies, but only because of their minimal market share.
10. <http://dataprivacyday2010.org/>.
11. Regulatory impacts on database administration are covered in depth in [Chapter 15](#), “[Regulatory Compliance and Database Administration](#).”
12. Even if they occur during regular working hours?
13. The META Group was acquired by Gartner, Inc., in 2004. If the dated META Group research note is still available, it would have to be acquired through Gartner. The META Group research note was published March 20, 1998 (Open Computing & Server Strategies, File: 656).

Chapter 2

1. If you prefer commercial software over open source, there are commercial offerings of some of the NoSQL products. For example, DataStax is based on Cassandra.
2. FileMaker is offered in a professional, multiuser version, too.
3. In DB2, the area used for caching program structures in memory is referred to as the EDM pool. In SQL Server it is called the SQL cache, and in Oracle two structures are used, the PGA and the shared pool in the SGA.
4. In DB2, system parameters are set by assembling the DSNZPARM member. SQL Server uses the SP_CONFIGURE system procedure to set system parameters, and

Oracle parameters are controlled using INIT.ORA.

5. In DB2, the SQL interface is referred to as SPUFI. IBM also provides Data Studio for GUI-based SQL creation and submission. SQL Server calls the interface ISQL, and when using Oracle you can choose to submit SQL using SQL*Plus or the SQL Worksheet in Oracle Enterprise Manager.
6. Be careful, too, to examine the specifications for any new DBMS version or release. Sometimes features and functionality are removed from the DBMS, which might result in having to spend additional money to replace the lost functionality. For example, IBM removed its formerly free database utilities from DB2 between Versions 6 and 7 and bundled them for sale.
7. When a feature is *deprecated* it is no longer supported in the software.

Chapter 3

1. Performance problems due to pessimistic locking concurrency are common when Microsoft Access developers try to upsize their single-user applications to Microsoft SQL Server or Oracle.

Chapter 4

1. Some DBMS products, such as DB2 for z/OS, offer hardware-assisted compression, which can mitigate the CPU overhead incurred when compressing and decompressing data.

Chapter 5

1. The Microsoft information about moving from OLE DB to ODBC can be found at <http://social.technet.microsoft.com/Forums/en/sqldataaccess/thread/e696d0ac-f8e2-4b19-8a08-7a357d3d780f>
2. There are options for deploying .NET applications on other operating systems, one of which is Mono (<http://www.go-mono.com/mono-downloads/download.html>).
3. For more details and specifics regarding XML, refer to the following Web sites: www.xml.org and www.w3.org/XML/.
4. For example, DB2 for z/OS.

Chapter 6

1. Realistically, most participants will have only a few of the traits described here. This means that the design review leader must be prepared to engage in the review process under less than ideal circumstances at times.

Chapter 7

1. It is possible that a physical change (for example, to add storage) will not need to be recorded in the logical and conceptual data models as certain physical characteristics of the database are not pertinent to those models. However, other changes, such as modifying a data type or adding a new column, will need to be reflected in the logical and conceptual data models.
2. You might choose to recreate the indexes earlier in the process here, for example, before the data is loaded. It should depend on whether the DBMS being used builds indexes more efficiently before data is in the table or after.

Chapter 8

1. Availability Groups in SQL Server 2012 offer a similar failover technology.
2. This type of clustering refers to hardware and differs from the clustering discussed in [Chapter 4](#), which refers to the organization of data on disk.
3. Known as LPARs (Logical PARTitions).

Chapter 9

1. Conducting extensive design reviews, as discussed in [Chapter 6](#), can help to ferret out the bad SQL before it hits production systems.
2. How can you know what *abnormal* is if you do not know what *normal* is?
3. Rick Sturm, Wayne Morris, and Mary Jander, *Foundations of Service Level Management* (Indianapolis, IN: SAMS Publishing, 2000).

Chapter 10

1. Or you have enabled bulk-logged recovery mode.
2. Oracle strongly advises customers to use cost-based optimization and plans to deprecate rule-based optimization in a future Oracle DBMS release.

Chapter 11

1. A multicore processor is a single computing component with two or more independent actual processors (called *cores*).
2. Raw partitions can be created in a Windows environment, too; it is just very rare.
3. Dynamic Management Views in Microsoft SQL Server are similar to Oracle Dynamic Performance Tables. DB2 for z/OS provides Real Time Statistics.

Chapter 14

1. This is a particularly vexing problem with purchased third-party applications. Many such applications rely on SA authority because it is easier. However, it is not a secure

practice.

Chapter 15

1. Thomas C. Redman, "Data: An Unfolding Quality Disaster," *DMReview*, August 2004.
2. <http://hollistibbetts.sys-con.com/node/1975126>.
3. Thomas C. Redman, *Data Quality: Management and Technology* (New York, NY: Bantam, 1992).
4. www.computerworld.com/s/article/9026166/Database_admin_steals_2.3M_consumer
5. In 2007, the SNIA (Storage Networking Industry Association) Data Forum published the 100 Year Archive Requirements Survey. The survey validates the need for long-term data retention. Eighty percent of the survey's respondents declared they have information they must keep more than 50 years, and 68 percent of respondents said they must keep it more than 100 years.
www.sresearch.com/100_Year_Archive_Requirements_Survey_and_Report.html.

Chapter 16

1. Some recovery options or products can find quiet points on the database log without requiring quiesce during image copy backups.
2. By default, this occurs about once every minute.

Chapter 17

1. Not including damage to roads and bridges or public buildings, as well as contents inside buildings and residences.

Chapter 18

1. DBAs may work in conjunction with a storage administration group, if one exists at their shop.
2. www.informationweek.com/816/gerstner.htm.
3. www.emc.com/collateral/demos/microsites/emc-digital-universe-2011/index.htm.
4. "The Petabyte Challenge: 2011 IOUG Database Growth Survey," conducted by Unisphere Research and sponsored by Oracle.
5. www.wintercorp.com/WhitePapers/WC_TopTenWP.pdf.
6. Lou Agosta, "How Much Data Is on the Planet?" *Giga IdeaByte*, September 20, 2000.
7. Either tools supplied by the DBMS vendor or purchased from a third-party vendor.
8. An interesting article on the topic can be found online at

www.teradatamagazine.com/v11n03/Tech2Tech/Why-Multi-Temperature-Data-Matters/.

Chapter 19

1. Loading before indexes are built is not a good idea for Sybase and Microsoft SQL Server because clustered indexes directly affect the physical ordering of rows in the table.
2. In most cases the CPU increase should be minimal. The additional CPU is required to combine the results from the parallel operations back into a single result.

Chapter 20

1. Ralph Kimball, *The Data Warehouse Toolkit* (New York, NY: John Wiley & Sons, 1996).

Chapter 22

1. www.pbs.org/wgbh/roadshow/.

Chapter 23

1. In most cases, the DBMS vendor DBA tools support only the DBMS products offered by that vendor. For example, you would generally not expect Oracle DBA tools to manage Microsoft SQL Server databases.

Chapter 24

1. IBMLink is a proprietary network that IBM opens up only to its customers.

```
CREATE TABLE ITEM
  (ItemNum      integer      not null,
   ItemSize     CHAR(1),
   ItemColor    CHAR(10),
   ItemDescr    CHAR(100)
);
```



```
CREATE TABLE ITEM
  (ItemNum      INTEGER      NOT NULL,
   ItemSize    CHAR(1),
   ItemColor   CHAR(10),
   ItemDescr   CHAR(10)
);
CREATE TABLE ITEM_DESC
  (ItemNum      INTEGER      NOT NULL,
   ItemDesc     CHAR(90)
);
```

```
CREATE TABLE ITEM_DESC
  (ItemNum          INTEGER          NOT NULL,
   ItemCtr          INTEGER          NOT NULL,
   ItemDesc         CHAR(100))
;
```

```
CREATE TABLE CUST_BALANCE
  (CustNum      INTEGER      NOT NULL,
   BalancePeriod INTEGER     NOT NULL,
   Balance      DECIMAL(15,2),
   constraint PKCB PRIMARY KEY (CustNum, BalancePeriod)
  ;
```

```
CREATE TABLE CUST_BALANCE
(CustNum          INTEGER          NOT NULL,
 Period1_Balance DECIMAL(15,2),
 Period2_Balance DECIMAL(15,2),
 Period3_Balance DECIMAL(15,2),
 Period4_Balance DECIMAL(15,2),
 Period5_Balance DECIMAL(15,2),
 Period6_Balance DECIMAL(15,2),
 constraint PKCB PRIMARY KEY (CustNum)
;
```

```
CREATE TABLE DEPT
  (ParentDeptNum    INTEGER          NOT NULL,
   DeptName         CHAR(25),
   SupervisorNum    INTEGER,
   ReportsToDeptNum INTEGER,
   constraint PKDN  PRIMARY KEY (DeptNum),
   constraint FKCB  FOREIGN KEY (ReportsToDeptNum) REFERENCES
DEPT
                                ON DELETE RESTRICT
;

```

```
CREATE TABLE DEPT
  (DeptNum      INTEGER      NOT NULL,
   ChildDeptNum INTEGER      NOT NULL,
   Level        INTEGER,
   Detail       CHAR(1),
   DeptName     CHAR(25),
   SupervisorNum INTEGER,
   constraint PKDN PRIMARY KEY (DeptNum, ChildDeptNum)
  );
```

```
SELECT CustName, PolicyNo, BenefitSummary
FROM InsurancePolicy
FOR BUSINESS_TIME AS OF '2012-04-15'
WHERE CustNo = ?
;
```

```
SELECT deptnum, deptname FROM dept WHERE supervisorum = '903';
```



```
SELECT empno, firstname, lastname
FROM employee
WHERE workdept IN
      (SELECT deptno
       FROM department
       WHERE deptname = 'DBA');
SELECT empno, firstname, lastname
FROM employee,
      department
WHERE workdept = deptno
AND deptname = 'DBA';
```

```
<CUSTOMER>  
  <first_name>Craig</first_name>  
  <middle_initial>S.</middle_initial>  
  <last_name>Mullins</last_name>  
  <company_name>Mullins Consulting, Inc.</company_name>  
  <street_address>15 Coventry Ct.</street_address>  
  <city>Sugar Land</city>  
  <state>TX</state>  
  <zip_code>77479</zip_code>  
  <country>USA</country>  
</CUSTOMER>
```

```
ALTER TABLE Table_1
  ADD COLUMN new_column INTEGER NULL
;
```

```
EXEC SP_DBOPTION 'pubs', 'trunc. log on chkpt.', 'false'
```

```
SELECT last_name, first_name, middle_initial, empno
FROM employee
WHERE position = 'MANAGER'
AND work_code = 1
AND dept = '001000';
```

```
SELECT last_name, first_name, middle_initial, empno
FROM employee
WHERE work_code = 1
AND dept = '001000';
```

```
SELECT last_name, first_name, middle_initial, empno, position
FROM employee
WHERE position in ('MANAGER', 'DIRECTOR', 'VICE PRESIDENT')
ORDER BY last_name;
```

```
WHERE column1 >= 1 AND column1 <= 100  
WHERE column1 BETWEEN 1 AND 100
```



```
SELECT    d.dept_name, e.last_name, e.empno
FROM      employee  e,
          department d
WHERE     e.deptno = d.deptno
AND      d.deptno = '808';
```

```
SELECT    d.dept_name, e.last_name, e.empno
FROM      employee  e,
          department d
WHERE     e.deptno = d.deptno
AND       e.deptno = '808';
```

```
EXPLAIN plan SET STATEMENT_ID = 'emptest' FOR
SELECT  position, last_name, first_name, middle_initial, empno
FROM    employee
WHERE   position IN ('MANAGER', 'DIRECTOR', 'VICE PRESIDENT')
ORDER BY position;
```

```
SELECT /*+ USE_NL */
       e.position, e.last_name, e.empno, d.manager
FROM   employee e,
       department d
WHERE  d.dept_id = e.dept_id
AND    position IN ('MANAGER', 'DIRECTOR', 'VICE PRESIDENT')
ORDER BY position;
```

```
SELECT last_name, first_name, empno, deptno
FROM employee
WHERE empno BETWEEN '001000' AND '009999'
AND (salary > 50000.00 OR 0 = 1)
ORDER by last_name;
```

```
SELECT    e.position, e.last_name, e.empno, d.manager
FROM      employee   e,
          department d
WHERE     d.dept_id = e.dept_id
AND       position = 'MANAGER'
OR        position = 'DIRECTOR'
OR        position = 'VICE PRESIDENT'
ORDER BY position;
```

```
SELECT  e.position, e.last_name, e.empno, d.manager
FROM    employee e,
        department d
WHERE   d.dept_id = e.dept_id
AND     position IN ('MANAGER', 'DIRECTOR', 'VICE PRESIDENT')
ORDER BY position;
```

```
SELECT    position, last_name, empno
FROM      employee
WHERE     dept_id LIKE '%X'
ORDER BY position;
```



```
SELECT    position, last_name, empno
FROM      employee
WHERE     dept_id LIKE 'A%X'
OR        dept_id LIKE 'B%X'
ORDER BY position;
```

```
CREATE TABLE EMP
(empno          INTEGER          PRIMARY KEY,
 emp_address    VARCHAR(70),
 emp_type       CHAR(8),
 emp_dept       CHAR(3)         NOT NULL WITH DEFAULT,
 salary         DECIMAL(7,2)    NOT NULL,
 commission    DECIMAL(7,2),
 bonus         DECIMAL(7,2)
) IN db.ts;
```

```
CREATE DISTINCT TYPE canadian_dollar AS DECIMAL(11,2);  
CREATE DISTINCT TYPE US_dollar AS DECIMAL(11,2);  
CREATE DISTINCT TYPE euro AS DECIMAL(11,2);  
CREATE DISTINCT TYPE japanese_yen AS DECIMAL(15,2);
```

```
CREATE TABLE EMP
(empno          INTEGER          PRIMARY KEY,
 CONSTRAINT check_empno
 CHECK (empno BETWEEN 100 and 25000),
emp_address    VARCHAR(70),
emp_type       CHAR(8)
 CHECK (emp_type IN ('temp', 'fulltime', 'contract')),
emp_dept       CHAR(3)          NOT NULL WITH DEFAULT,
salary         DECIMAL(7,2)     NOT NULL
 CONSTRAINT check_salary
 CHECK (salary < 50000.00),
commission     DECIMAL(7,2),
bonus          DECIMAL(7,2)
) IN db.ts;
```

```
CREATE TABLE EMP
(empno          INTEGER          PRIMARY KEY,
 CONSTRAINT check_empno
 CHECK (empno BETWEEN 100 AND 25000),
 emp_address    VARCHAR(70),
 emp_type       CHAR(8)
 CHECK (emp_type IN ('temp', 'fulltime', 'contract')),
 emp_dept       CHAR(3)          NOT NULL WITH DEFAULT,
 salary         DECIMAL(7,2)     NOT NULL
 CONSTRAINT check_salary
 CHECK (salary < 50000.00),
 commission     DECIMAL(7,2),
 bonus         DECIMAL(7,2)
 CONSTRAINT comm_vs_salary
 CHECK (salary > commission),
 CONSTRAINT comm_bonus
 CHECK (commission=0 OR bonus=0),
) IN db.ts;
```

```
CREATE rule state_rule AS @state IN ("IL", "WI", "IN")
EXEC sp_bindrule "state_rule", "publishers.state"
EXEC sp_bindrule "state_rule", "authors.state"
```

```
emp_type CHAR(8) DEFAULT 'fulltime'  
        CHECK (emp_type IN ('temp', 'fulltime', 'contract')), ...
```

```
emp_type CHAR(8) DEFAULT 'new'  
CHECK (emp_type IN ('temp', 'fulltime', 'contract')), ...
```



```
CREATE TRIGGER salary_update
  BEFORE UPDATE OF salary
  ON emp
  FOR EACH ROW MODE DB2SQL
  WHEN (new.salary > (old.salary * 1.5))
  BEGIN atomic
    SIGNAL SQLSTATE '75001' ('Raise exceeds 50%');
  END;
```

```
CREATE TRIGGER title_del
ON titles FOR DELETE
AS
IF @@rowcount = 0
    RETURN
DELETE titleauthor
    FROM titleauthor, deleted, title
    WHERE titles.title_id = deleted.title_id
RETURN
```

```
CREATE TRIGGER title_ins
ON titleauthor FOR INSERT
AS
DECLARE @rc int
SELECT @rc = @@rowcount
IF @rc = 0
    RETURN
IF (select count(*)
    FROM titles, inserted
    WHERE titles.title_id = inserted.title_id)!=@rc
    BEGIN
        RAISERROR 20001 "Invalid title: title_id
            does not exist on titles table"
        ROLLBACK transaction
        RETURN
    END
RETURN
```

```
CREATE TRIGGER title_upd
ON titles FOR UPDATE
AS
IF UPDATE (title_id)
  IF (select count(*)
      FROM deleted, titles
      WHERE deleted.title_id = title.titleid) = 0
  BEGIN
    UPDATE titleauthor
      SET titleauthor.titleid = null
    FROM titleauthor, deleted
    WHERE titleauthor.titleid = deleted.title_id
  END
RETURN
```

```
CREATE TABLE PROJECT
(projno      INTEGER      PRIMARY KEY,
deptno      INTEGER      NOT NULL,
mg_empno    INTEGER,
status      INTEGER      NOT NULL
            CHECK (status BETWEEN 0 AND 3),
start_date  DATE,
) IN db.ts;
```

```
CREATE TABLE COURSE
(COURSENO    INTEGER          NOT NULL,
 TITLE      VARCHAR(20)     NOT NULL,
 CREDITS     SMALLINT        NOT NULL WITH DEFAULT 3,
 PRICE      DECIMAL(7,2)    NOT NULL,
 BUS_START   DATE            NOT NULL,
 BUS_END     DATE            NOT NULL,

 PERIOD BUSINESS_TIME(BUS_START, BUS_END),

 PRIMARY KEY(COURSENO, BUSINESS_TIME WITHOUT OVERLAPS)
);
```

```
SELECT Price
FROM Course FOR BUSINESS_TIME AS OF '2012-04-03'
WHERE Title = 'ENGLISH 101';
```

```
UPDATE Course
FOR PORTION OF BUSINESS_TIME FROM '2011-11-01' TO '2012-02-02'
SET Price = 325.00
WHERE CourseNo = 100;
```



```
CREATE role MANAGER;  
COMMIT;  
GRANT select, insert, update, delete on employee to MANAGER;  
GRANT select, insert, update, delete on job_title to MANAGER;  
GRANT execute on payroll to MANAGER;  
COMMIT;  
GRANT MANAGER to user1;  
COMMIT;
```

```
CREATE view emp_all
AS
SELECT first_name, last_name, middle_initial,
       street_address, state, zip_code
FROM   employee;
```

```
CREATE view emp_dept20
AS
SELECT first_name, last_name, middle_initial,
       street_address, state, zip_code
FROM   employee
WHERE  deptno = 20;
```

```
statement := "SELECT * FROM userinfo WHERE id = " + in_var + ";"
```

```
SELECT "CREATE DATABASE ", dbname, " BUFFERPOOL ", bpname,  
      " STOGROUP ", sname, " ;"  
FROM   sysdatabases;
```

LOAD . . . STATUS POSITION (20:25) CHAR(6) NULLIF(26)='N' . . .