

Berkeley Db Java Edition Tutorial

Written by experienced Oracle insiders, this essential guide distills a vast amount of information into an easy-to-read volume that covers every aspect of the Oracle database. Readers of all technical levels will learn about Oracle's features and technologies, including the product line, architecture, data structures, networking, concurrency, tuning and much more. Augmented with illustrations and helpful hints, the fifth edition of Oracle Essentials offers a valuable one-stop overview of Oracle Database 12c, Oracle's newest database release. More comprehensible than huge complete references, and more detailed than most primers, this book gives current Oracle users the conceptual background they need to understand how the Oracle database truly works. For those new to Oracle, this all-in-one guide provides an essential introduction that will get them up to speed. Joe Celko's Complete Guide to NoSQL provides a complete overview of non-relational technologies so that you can become more nimble to meet the needs of your organization. As data continues to explode and grow more complex, SQL is becoming less useful for querying data and extracting meaning. In this new world of bigger and faster data, you will need to leverage non-relational technologies to get the most out of the information you have. Learn where, when, and why the benefits of NoSQL outweigh those of SQL with Joe Celko's Complete Guide to NoSQL. This book covers three areas that make today's new data different from the data of the past: velocity, volume and variety. When information is changing faster than you can collect and query it, it simply cannot be treated the same as static data. Celko will help you understand velocity, to equip you with the tools to drink from a fire hose. Old storage and access models do not work for big data. Celko will help you understand volume, as well as different ways to store and access data such as petabytes and exabytes. Not all data can fit into a relational model, including genetic data, semantic data, and data generated by social networks. Celko will help you understand variety, as well as the alternative storage, query, and management frameworks needed by certain kinds of data. Gain a complete understanding of the situations in which SQL has more drawbacks than benefits so that you can better determine when to utilize NoSQL technologies for maximum benefit Recognize the pros and cons of columnar, streaming, and graph databases Make the transition to NoSQL with the expert guidance of best-selling SQL expert Joe Celko

The design and analysis of efficient data structures has long been recognized as a key component of the Computer Science curriculum. Goodrich, Tomassia and Goldwasser's approach to this classic topic is based on the object-oriented paradigm as the framework of choice for the design of data structures. For each ADT presented in the text, the authors provide an associated Java interface. Concrete data structures realizing the ADTs are provided as Java classes implementing the interfaces. The Java code implementing fundamental data structures in this book is organized in a single Java package, net.datastructures. This package forms a coherent library of data structures and algorithms in Java specifically designed for educational purposes in a way that is complimentary with the Java Collections Framework.

Python Essential Reference, 3rd Edition, is a comprehensive reference to the Python programming language. The focus of this latest edition is to add coverage of significant new features and new library modules added to the language over the past five years. Clearly written with concise organization, the new features covered include new style classes, unification of types and classes, xmlrpclip, intertools, bz2 and optparse, making it the most up-to-date Python book on the market.

Explains how to access and create MySQL databases through PHP scripting, including discussion of authentication, network connectivity, sessions, and content management.

Database Administration

Semantics in Business Systems

Version Control with Subversion

The Savvy Manager's Guide

Python Essential Reference

The Java EE 7 Tutorial: Volume 2, Fifth Edition, is a task-oriented, example-driven guide to developing enterprise applications for the Java Platform, Enterprise Edition 7 (Java EE 7). Written by members of the Java EE documentation team at Oracle, this book provides new and intermediate Java programmers with a deep understanding of the platform. This guide includes descriptions of platform features and provides instructions for using the latest versions of NetBeans IDE and GlassFish Server Open Source Edition. The book introduces Enterprise JavaBeans components, the Java Persistence API, the Java Message Service (JMS) API, Java EE security, transactions, resource adapters, Java EE Interceptors, Batch Applications for the Java Platform, and Concurrency Utilities for Java EE. The book culminates with three case studies that illustrate the use of multiple Java EE 7 APIs.

RPM is the Linux industry standard for making application installation easy and its use is gaining mindshare amongst users and administrators of other Unix platforms Offers insight and examples to creating applications that rely upon or enhance RPM, enabling users to package and deploy software in RPM format Discusses the use of RPM to manage software and examines the tools provided for user control Book will be technically reviewed by the key RPM programmer at Red Hat RED HAT PRESS(TM) Linux Solutions from the Experts at Red Hat Red Hat-the world's leading Linux company-presents a series of unrivaled guides that are reviewed and approved by the experts at Red Hat. Each book is packed with invaluable tips and techniques that are ideal for everyone from beginning to advanced network and systems professionals, as well as home and small businesses.

This is the first book to devote complete coverage to the most recent release of the popular embedded open source database SQLite. The book acts as both an ideal tutorial and reference guide. It offers experienced database developers a thorough overview of its capabilities and APIs, yet is mindful of newcomers who may be making their first foray into the database environment with SQLite. Readers are presented with introductions to the SQLite extensions available for C, Java, Perl, PHP, Python, Ruby, and Tcl.

This book constitutes the thoroughly refereed postproceedings of the 3rd and 4th International Workshop on Databases, Information Systems and Peer-to-Peer Computing, DBISP2P 2005 and DBISP2P 2006, held in Trondheim, Norway, in August 2005 and in Seoul, Korea, in September 2006, as satellite events of VLDB, the International Conference on Very Large Data Bases.

Outside of the world of enterprise computing, there is one database that enables a huge range of software and hardware to flex relational database capabilities, without the baggage and cost of traditional database management systems. That database is SQLite—an embeddable database with an amazingly small footprint, yet able to handle databases of enormous size. SQLite comes equipped with an array of powerful features available through a host of programming and development environments. It is supported by languages such as C, Java, Perl, PHP, Python, Ruby, TCL, and more. The Definitive Guide to SQLite, Second Edition is devoted to complete coverage of the latest version of this powerful database. It offers a thorough overview of SQLite’s capabilities and APIs. The book also uses SQLite as the basis for helping newcomers make their first foray into database development. In only a short time you can be writing programs as diverse as a server-side browser plug-in or the next great iPhone or Android application! Learn about SQLite extensions for C, Java, Perl, PHP, Python, Ruby, and Tcl. Get solid coverage of SQLite internals. Explore developing iOS (iPhone) and Android applications with SQLite. SQLite is the solution chosen for thousands of products around the world, from mobile phones and GPS devices to set-top boxes and web browsers. You almost certainly use SQLite every day without even realizing it!

Java, A Beginner's Guide, 5th Edition

The Best Guide to Database Programming with Java GUI, PostgreSQL, and SQL Server

Subversion 1.6 Official Guide

Data Management. Data, Data Everywhere

International Workshops, DBISP2P 2005/2006, Trondheim, Norway, August 28-29, 2006, Revised Selected Papers

This book offers the straightforward, practical answers you need to help you do your job. This hands-on tutorial/reference/guide to PostgreSQL and SQL Server is not only perfect for students and beginners, but it also works for experienced developers who aren't getting the most from PostgreSQL and SQL Server. As you would expect, this book shows how to build from scratch two different databases: PostgreSQL and SQL Server using Java. In designing a GUI and as an IDE, you will make use of the NetBeans tool. In chapter one, you will learn: How to install NetBeans, JDK 11, and the PostgreSQL connector; How to integrate external libraries into projects; How the basic PostgreSQL commands are used; How to query statements to create databases, create tables, fill tables, and manipulate table contents is done. In chapter two, you will learn querying data from the postgresql using jdbc including establishing a database connection, creating a statement object, executing the query, processing the resultset object, querying data using a statement that returns multiple rows, querying data using a statement that has parameters, inserting data into a table using jdbc, updating data in postgresql database using jdbc, calling postgresql stored function using jdbc, deleting data from a postgresql table using jdbc, and postgresql jdbc transaction. In chapter three, you will learn the basics of cryptography using Java. Here, you will learn how to write a Java program to count Hash, MAC (Message Authentication Code), store keys in a KeyStore, generate PrivateKey and PublicKey, encrypt / decrypt data, and generate and verify digital prints. You will also learn how to create and store salt passwords and verify them. In chapter four, you will create a PostgreSQL database, named Bank, and its tables. In chapter five, you will create a Login table. In this case, you will see how to create a Java GUI using NetBeans to implement it. In addition to the Login table, in this chapter you will also create a Client table. In the case of the Client table, you will learn how to generate and save public and private keys into a database. You will also learn how to encrypt / decrypt data and save the results into a database. In chapter six, you will create an Account table. This account table has the following ten fields: account_id (primary key), client_id (primarykey), account_number, account_date, account_type, plain_balance, cipher_balance, decipher_balance, digital_signature, and signature_verification. In this case, you will learn how to implement generating and verifying digital prints and storing the results into a database. In chapter seven, you create a table named Client_Data, which has seven columns: client_data_id (primary key), account_id (primary_key), birth_date, address, mother_name, telephone, and photo_path. In chapter eight, you will be taught how to create a SQL Server database, named Crime, and its tables. In chapter nine, you will be taught how to extract image features, utilizing BufferedImage class, in Java GUI. In chapter ten, you will be taught to create Java GUI to view, edit, insert, and delete Suspect table data. This table has eleven columns: suspect_id (primary key), suspect_name, birth_date, case_date, report_date, suspect_status, arrest_date, mother_name, address, telephone, and photo. In chapter eleven, you will be taught to create Java GUI to view, edit, insert, and delete Feature_Extraction table data. This table has eight columns: feature_id (primary key), suspect_id (foreign key), feature1, feature2, feature3, feature4, feature5, and feature6. In chapter twelve, you will add two

tables: *Police_Station* and *Investigator*. These two tables will later be joined to *Suspect* table through another table, *File_Case*, which will be built in the seventh chapter. The *Police_Station* has six columns: *police_station_id* (primary key), *location*, *city*, *province*, *telephone*, and *photo*. The *Investigator* has eight columns: *investigator_id* (primary key), *investigator_name*, *rank*, *birth_date*, *gender*, *address*, *telephone*, and *photo*. Here, you will design a Java GUI to display, edit, fill, and delete data in both tables. In chapter thirteen, you will add two tables: *Victim* and *File_Case*. The *File_Case* table will connect four other tables: *Suspect*, *Police_Station*, *Investigator* and *Victim*. The *Victim* table has nine columns: *victim_id* (primary key), *victim_name*, *crime_type*, *birth_date*, *crime_date*, *gender*, *address*, *telephone*, and *photo*. The *File_Case* has seven columns: *file_case_id* (primary key), *suspect_id* (foreign key), *police_station_id* (foreign key), *investigator_id* (foreign key), *victim_id* (foreign key), *status*, and *description*. Here, you will also design a Java GUI to display, edit, fill, and delete data in both tables. Finally, this book is hopefully useful and can improve database programming skills for every Java/PostgreSQL/SQL Server programmer.

This is the official guide and reference manual for Subversion 1.6 - the popular open source revision control technology.

InfoWorld is targeted to Senior IT professionals. Content is segmented into Channels and Topic Centers. InfoWorld also celebrates people, companies, and projects.

This is the first book to cover db4o programming in comprehensive detail. Readers are briefed on all of the topics necessary to begin using it in production environments, including installation and configuration, querying and managing objects, performing transactions, and data replication. Newcomers to the topic aren't forgotten, as early chapters are devoted to object database fundamentals, in addition to technical considerations and migration strategies. Complete with numerous C# and Java examples, readers will be able to follow along with the examples regardless of their chosen language.

The Java EE 7 Tutorial: Volume 1, Fifth Edition, is a task-oriented, example-driven guide to developing enterprise applications for the Java Platform, Enterprise Edition 7 (Java EE 7). Written by members of the Java EE documentation team at Oracle, this book provides new and intermediate Java programmers with a deep understanding of the platform. This guide includes descriptions of platform features and provides instructions for using the latest versions of NetBeans IDE and GlassFish Server Open Source Edition. The book introduces platform basics, including resource creation, resource injection, and packaging. It covers JavaServer Faces, Java Servlets, the Java API for WebSocket, the Java API for JSON Processing (JSON-P), internationalization and localization, Bean Validation, Contexts and Dependency Injection for Java EE (CDI), and web services (JAX-WS and JAX-RS).

Professional NoSQL

A Cyber-Physical Systems Approach

The Definitive Guide to db4o

Third International Conference, CNSA 2010, Chennai, India, July 23-25, 2010 Proceedings

The Java EE 6 Tutorial

Teaches a model for robust programming and code reuse, covering the standard JDBC and RMI Java packages for enhanced communication with remote databases and improved object-oriented programming. Original. (Advanced).

While standardization has empowered the software industry to substantially scale software development and to provide affordable software to a broad market, it often does not address smaller market segments, nor the needs and wishes of individual customers. Software product lines reconcile mass production and standardization with mass customization in software engineering. Ideally, based on a set of reusable parts, a software manufacturer can generate a software product based on the requirements of its customer. The concept of features is central to achieving this level of automation, because features bridge the gap between the requirements the customer has and the functionality a product provides. Thus features are a central concept in all phases of product-line development. The authors take a developer's viewpoint, focus on the development, maintenance, and implementation of product-line variability, and especially concentrate on automated product derivation based on a user's feature selection. The book consists of three parts. Part I provides a general introduction to feature-oriented software product lines, describing the product-line approach and introducing the product-line development process with its two elements of domain and application engineering. The pivotal part II covers a wide variety of implementation techniques including design patterns, frameworks, components, feature-oriented programming, and aspect-oriented

programming, as well as tool-based approaches including preprocessors, build systems, version-control systems, and virtual separation of concerns. Finally, part III is devoted to advanced topics related to feature-oriented product lines like refactoring, feature interaction, and analysis tools specific to product lines. In addition, an appendix lists various helpful tools for software product-line development, along with a description of how they relate to the topics covered in this book. To tie the book together, the authors use two running examples that are well documented in the product-line literature: data management for embedded systems, and variations of graph data structures. They start every chapter by explicitly stating the respective learning goals and finish it with a set of exercises; additional teaching material is also available online. All these features make the book ideally suited for teaching - both for academic classes and for professionals interested in self-study.

The latest edition of a popular text and reference on database research, with substantial new material and revision; covers classical literature and recent hot topics. Lessons from database research have been applied in academic fields ranging from bioinformatics to next-generation Internet architecture and in industrial uses including Web-based e-commerce and search engines. The core ideas in the field have become increasingly influential. This text provides both students and professionals with a grounding in database research and a technical context for understanding recent innovations in the field. The readings included treat the most important issues in the database area--the basic material for any DBMS professional. This fourth edition has been substantially updated and revised, with 21 of the 48 papers new to the edition, four of them published for the first time. Many of the sections have been newly organized, and each section includes a new or substantially revised introduction that discusses the context, motivation, and controversies in a particular area, placing it in the broader perspective of database research. Two introductory articles, never before published, provide an organized, current introduction to basic knowledge of the field; one discusses the history of data models and query languages and the other offers an architectural overview of a database system. The remaining articles range from the classical literature on database research to treatments of current hot topics, including a paper on search engine architecture and a paper on application servers, both written expressly for this edition. The result is a collection of papers that are seminal and also accessible to a reader who has a basic familiarity with database systems.

The Java EE 6 Tutorial: Advanced Topics, Fourth Edition, is a task-oriented, example-driven guide to developing enterprise applications for the Java Platform, Enterprise Edition 6 (Java EE 6). Written by members of the Java EE 6 documentation team at Oracle, this book provides new and intermediate Java programmers with a deep understanding of the platform. This guide-which builds on the concepts introduced in The Java EE 6 Tutorial: Basic Concepts, Fourth Edition-contains advanced material, including detailed introductions to more complex platform features and instructions for using the latest version of the NetBeans IDE and the GlassFish Server, Open Source Edition. This book introduces the Java Message Service (JMS) API and Java EE Interceptors. It also describes advanced features of JavaServer Faces, Servlets, JAX-RS, Enterprise JavaBeans components, the Java Persistence API, Contexts and Dependency Injection for the Java EE Platform, web and enterprise application security, and Bean Validation. The book culminates with three new case studies that illustrate the use of multiple Java EE 6 APIs.

This book teaches the basics of XML with an original approach, using real-world examples from an interesting (and operating) environment with broad applicability. It covers the full spectrum of Berkeley DB XML tools, including the command-line shell, transactions, rollbacks, replication, archiving and monitoring. Techniques and concepts that have broad applicability outside of the subject matter are skillfully explained: XML, XPath, XQuery, XML schemas, all industry-standard technologies that find one of their best tutorial treatments, and all in the context of a simple database solution. The book also presents a remarkable example of query power.

24th British National Conference on Databases, BNCOD 24, Glasgow, UK, July 3-5, 2007, Proceedings

The Complete Guide to DBA Practices and Procedures

Oracle Essentials

The Definitive Guide to MySQL

Databases, Information Systems, and Peer-to-Peer Computing

This book provides experienced Access users who are novice programmers with frequently overlooked concepts and techniques necessary to create effective database applications. It focuses on designing effective tables in a multi-table application; using the Access interface or Access SQL to construct queries; and programming using the Data Access Object (DAO) and Microsoft Access object

models.

The Definitive Guide to Berkeley DB XML covers Sleepycat's Berkeley DB XML, an open source embedded XML database. Berkeley DB XML runs on all major operating systems and has support for the most popular programming languages. The book includes tutorials and complete language references for C++, Java, Perl, Python, and PHP. Berkeley DB XML has the potential to dramatically simplify the development of your application's data tier. With native XML storage, flexible indexing, and the powerful XQuery query language (think SQL for XML), BDB XML provides everything you need for efficient XML management. Combined with Berkeley DB's transactions, logging, and replication, BDB XML is a powerful document storage solution. Author Danny Brian presents XML basics from the ground up for XML beginners, and covers database installation and configuration, the core concepts behind Berkeley DB XML, a look at application architecture, demonstrations of simple operations (i.e., text queries) to complex (i.e., set computations) using XQuery and FLWOR expressions, document indexing options, and a thorough look at the database APIs. You ' ll learn tips and tricks that you can immediately apply to your own applications. Table of Contents A Quick Look at Berkeley DB XML The Power of an Embedded XML Database Installation and Configuration Getting Started Environments, Containers, and Documents Indexes XQuery with BDB XML BDB XML with C++ BDB XML with Python BDB XML with Java BDB XML with Perl BDB XML with PHP Managing Databases

Essential Skills--Made Easy! Learn the fundamentals of Java programming in no time from bestselling programming author Herb Schildt. Fully updated to cover Java Platform, Standard Edition 7 (Java SE 7), Java: A Beginner's Guide, Fifth Edition starts with the basics, such as how to compile and run a Java program, and then discusses the keywords, syntax, and constructs that form the core of the Java language. You'll also find coverage of some of Java's most advanced features, including multithreaded programming and generics. An introduction to Swing concludes the book. Get started programming in Java right away with help from this fast-paced tutorial. Designed for Easy Learning: Key Skills & Concepts--Chapter-opening lists of specific skills covered in the chapter Ask the Expert--Q&A sections filled with bonus information and helpful tips Try This--Hands-on exercises that show you how to apply your skills Self Tests--End-of-chapter questions that test your understanding Annotated Syntax--Example code with commentary that describes the programming techniques being illustrated

This step-by-step guide to explore database programming using Java is ideal for people with little or no programming experience. The goal of this concise book is not just to teach you Java, but to help you think like a programmer. Each brief chapter covers the material for one week of a college course to help you practice what you've learned. As you would expect, this book shows how to build from scratch two different databases: PostgreSQL and SQLite using Java. In designing a GUI and as an IDE, you will make use of the NetBeans tool. In the first chapter, you will learn: How to install NetBeans, JDK 11, and the PostgreSQL connector; How to integrate external libraries into projects; How the basic PostgreSQL commands are used; How to query statements to create databases, create tables, fill tables, and manipulate table contents is done. In the first chapter, you will learn: How to install NetBeans, JDK 11, and the PostgreSQL connector; How to integrate external libraries into projects; How the basic PostgreSQL commands are used; How to query statements to create databases, create tables, fill tables, and manipulate table contents is done. In the second chapter, you will learn querying data from the postgresql using jdbc including establishing a database connection, creating a statement object, executing the query, processing the resultset object, querying data using a statement that returns multiple rows, querying data using a statement that has parameters, inserting data into a table using jdbc, updating data in postgresql database using jdbc, calling postgresql stored function using jdbc, deleting data from a postgresql table using jdbc, and postgresql jdbc transaction. In chapter three, you will create a PostgreSQL database, named School, and its tables. In chapter four, you will study: Creating the initial three table projects in the school database: Teacher table, TClass table, and Subject table; Creating database configuration files; Creating a Java GUI for viewing and navigating the contents of each table; Creating a Java GUI for inserting and editing tables; and Creating a Java GUI to join and query the three tables. In chapter five, you will learn: Creating the main form to connect all forms; Creating a project will add three more tables to the school database: the Student table, the Parent table, and Tuition table; Creating a Java GUI to view and navigate the contents of each table; Creating a Java GUI for editing, inserting, and deleting records in each table; Creating a Java GUI to join and query the three tables and all six. In chapter six, you will study how to query the six tables. In chapter seven, you will be shown how to create SQLite database and tables with Java. In chapter eight, you will be taught how to extract image features, utilizing BufferedImage class, in Java GUI. Digital image techniques to extract image features used in this chapted are grascaling, sharpening, invertering, blurring, dilation, erosion, closing, opening, vertical prewitt, horizontal prewitt, Laplacian, horizontal sobel, and vertical sobel. For readers, you can develop it to store other advanced image features based on descriptors such as SIFT and others for developing descriptor based matching. In chapter nine, you will be taught to create Java GUI to view, edit, insert, and delete Suspect table data. This table has eleven columns: suspect_id (primary key), suspect_name, birth_date, case_date, report_date, suspect_status, arrest_date, mother_name, address, telephone, and photo. In chapter ten, you will be taught to create Java GUI to view, edit, insert, and delete Feature_Extraction table data. This table has eight columns: feature_id (primary key), suspect_id (foreign key), feature1, feature2, feature3, feature4, feature5, and feature6. All six fields (except keys) will have a BLOB data type, so that the image of the feature will be directly saved into this table. In chapter eleven, you will add two tables: Police_Station and Investigator. These two tables will later be joined to Suspect table through another table, File_Case, which will be built in the seventh chapter. The Police_Station has six columns: police_station_id (primary key), location, city, province, telephone, and photo. The Investigator has eight columns: investigator_id (primary key), investigator_name, rank, birth_date, gender, address, telephone, and photo. Here, you will design a Java GUI to display, edit, fill, and delete data in both tables. In chapter twelve, you will add two tables: Victim and Case_File. The File_Case table will connect four other tables: Suspect, Police_Station, Investigator and Victim. The Victim table has nine columns: victim_id (primary key), victim_name, crime_type, birth_date, crime_date, gender, address, telephone, and photo. The Case_File has seven columns: case_file_id (primary key), suspect_id (foreign key), police_station_id (foreign key), investigator_id (foreign key), victim_id (foreign key), status, and description. Here, you will also design a Java GUI to display, edit, fill, and delete data in both tables. Finally, this book is hopefully useful and can improve database programming skills for every Java/PostgreSQL/SQLite pogrammer.

This book features the refereed proceedings from the 24th British National Conference on Databases, held in Glasgow, Scotland in July 2007. The eighteen full papers and seven poster papers are presented, together with two invited contributions. Papers are organized into topical sections covering data applications, searching XML documents, querying XML documents, XML transformation, clustering and security, data mining, and extraction.

The CIO's Guide to Oracle Products and Solutions

The Definitive Guide to Berkeley DB XML

Introduction to Embedded Systems

Access Database Design & Programming

Red HatRPM Guide

Master Oracle NoSQL Database Enable highly reliable, scalable, and available data. Oracle NoSQL Database: Real-Time Big Data Management for the Enterprise shows you how to take full advantage of this cost-effective solution for storing, retrieving, and updating high-volume, unstructured data. The book covers installation, configuration, application development, capacity planning and sizing, and integration with other enterprise data center products. Real-world examples illustrate the concepts presented in this Oracle Press guide. Understand Oracle NoSQL Database architecture and the underlying data storage engine, Oracle Berkeley DB Install and configure Oracle NoSQL Database for optimal performance Develop complex, distributed applications using a rich set of APIs Read and write data into the Oracle NoSQL Database key-value store Apply an Avro schema to the value portion of the key-value pair using Avro bindings Learn best practices for capacity planning and sizing an enterpriselevel Oracle NoSQL Database deployment Integrate Oracle NoSQL Database with Oracle Database, Oracle Event Processing, and Hadoop Code examples from the book are available for download at www.OraclePressBooks.com.

** Expanded and revised in light of the GNU Compiler Collection (GCC) 4 release in April 2005, this book offers detailed coverage of GCC's somewhat daunting array of options and features and includes several chapters devoted to its support for languages like C, C++, Java, Objective-C, and Fortran. * Though targeting beginner and intermediate developers, this book goes well beyond basic compiler usage, combining instruction of GCC's advanced features and utilities (authconf, libtool, and gprof) with key coding techniques, such as profiling and optimization to show how to build and manage enterprise-level applications. * This is an enormous market. GCC is the defacto compiler collection for hundreds of thousands of open source projects worldwide, a wide variety of commercial development projects, and is the standard compiler for academic programs. From operating systems to the cloud, Oracle's products and services are everywhere, and it has the market share to prove it. Given the share diversity of the Oracle product line, and the level of complexity of integration, management can be quite a daunting task. The CIO's Guide to Oracle Products and Solutions is the go-to guide for all things Oracle. It provides management-level guidance on how to successfully navigate and manage the full range of Oracle products. The book presents management best practices and user/developer lessons learned in the use of Oracle products and services. Supplying both conceptual and technical views, the text focuses on what CIOs need to do to orient, or reorient, their organization toward the use of Oracle products and services. It describes how to develop a strategic framework for the use of these products and services rather than the specific product or service itself. This strategic framework will help you to prepare, educate, keep up with change, mitigate risk, and implement with the confidence needed to succeed. Providing an overview of the suite of Oracle technologies and solutions, the book covers the heart of the Oracle products set, including Oracle analytics, enterprise performance management, Oracle cloud, data management, application development, social business, and fusion. It examines compliance and security issues and includes metrics to help you evaluate potential solutions. The book also provides readers with access to a set of helpful resources on the book's page at www.crcpress.com, including cloud procurement best practices, cloud migration tips, a sample project procurement plan template, and various glossaries.*

A hands-on guide to leveraging NoSQL databases NoSQL databases are an efficient and powerful tool for storing and manipulating vast quantities of data. Most NoSQL databases scale well as data grows. In addition, they are often malleable and flexible enough to accommodate semi-structured and sparse data sets. This comprehensive hands-on guide presents fundamental concepts and practical solutions for getting you ready to use NoSQL databases. Expert author Shashank Tiwari begins with a helpful introduction on the subject of NoSQL, explains its characteristics and typical uses, and looks at where it fits in the application stack. Unique insights help you choose which NoSQL solutions are best for solving your specific data storage needs. Professional NoSQL: Demystifies the concepts that relate to NoSQL databases, including column-family oriented stores, key/value databases, and document databases. Delves into installing and configuring a number of NoSQL products and the Hadoop family of products. Explains ways of storing, accessing, and querying data in NoSQL databases through examples that use MongoDB, HBase, Cassandra, Redis, CouchDB, Google App Engine Datastore and more. Looks at architecture and internals. Provides guidelines for optimal usage, performance tuning, and scalable configurations. Presents a number of tools and utilities relating to NoSQL, distributed platforms, and scalable processing, including Hive, Pig, RRDtool, Nagios, and more.

The CIO's Guide to Oracle Products and SolutionsCRC Press

MongoDB in Action

What Every SQL Professional Needs to Know about Non-Relational Databases

Java In Action: An Excellent Guide to Explore JDBC And Database Applications

Feature-Oriented Software Product Lines

Advanced Environments, Tools, and Applications for Cluster Computing

The book illustrates how this applies to the future of application system development, especially how it informs and affects business rule-based approaches, and how semantics will play out with XML and the semantic Web. The book also contains a reference guide to related terms and technologies.

Besides covering the most recently released versions of GCC, this book provides a complete command reference, explains how to use the online help system, and covers material not covered in other texts, including profiling, test coverage, and how to build applications on a variety of operating system and hardware platforms. It also covers how to integrate with other GNU development tools like automake, autoconf, and libtool.

The Third International Conference on Network Security and Applications (CNSA-2010) focused on all technical and practical aspects of network security and its applications for wired and wireless networks. The goal of this conference is to bring together researchers and practitioners from academia and industry to focus on understanding modern security threats and countermeasures, and establishing new research directions in these areas. Authors are invited to contribute to the conference by submitting articles that illustrate research results, practical work and industrial experiences describing significant advances in the areas of security and its applications, including:

- Network Security
- Wireless Network Security
- Mobile, Ad Hoc and Sensor Network Security
- Peer-to-Peer Network Security
- Database and System Security
- Intrusion Detection and Prevention
- Internet Security, and Applications Security and Network Management
- E-mail Security

Phishing, E-mail Fraud • Virus, Worms, Trojon Protection • Security Threats and Countermeasures (DDoS, MiM, Session Hijack, Replay attack etc.) • Ubiquitous Computing Security • Web 2. 0 Security • Cryptographic Protocols • Performance Evaluation Protocols and Security Application There were 182 submissions to the conference and the Program Committee selected 63 publication. The book is organized as a collection of papers from the First International Workshop on Trust Management in P (IWTMP2PS 2010), the First International Workshop on Database Management Systems (DMS- 2010), and the First Internat Workshop on Mobile, Wireless and Networks Security (MWNS-2010).

The Berkeley DB Book is intended to be a practical guide to the intricacies of Berkeley DB; an in-depth analysis of the complex issues which are often covered in terse footnotes in the dense Berkeley DB reference manual. It explains the technology at a and also covers the internals with generous code and design examples. Berkeley DB is becoming the database of choice for a makers and for in memory cache of large scale applications like search engines and high traffic web sites.

Small, special-purpose computing devices and high-end core Internet servers need fast, reliable database management. Berkeley embedded database that provides high-performance, scalable, transaction-protected and recoverable data management service applications. Extremely portable, this library runs under almost all UNIX and Windows variants, as well as a number of embedded time operating systems. Berkeley DB is the ultimate resource for the world's most widely deployed embedded database engine aid software architects and engineers, product managers, and systems and network administrators without the overhead im database products. Designed by programmers for programmers, this classic library style toolkit provides a broad base of func application writers. This book will help you to make intelligent choices about when and how to use Berkeley DB to meet your can visit the Sleepycat website to get the latest errata for this book. NOTE: The first printing of this book contained an erro contents that caused the page numbers to be off. This will be corrected in the second printing. If you have an earlier edition download a pdf of the correct table of contents that you can print out and use with your book. If you have any questions, p contact the editor of this book at stephanie.wall@newriders.com.

Real-Time Big Data Management for the Enterprise

PHP and MySQL for Dynamic Web Sites

Berkeley DB

Database Programming with JDBC and Java

Covers MongoDB version 3.0

Database Administration, Second Edition , is the definitive, technology-independent guide to the modern discipline of database administration. Packed with best practices and proven solutions for any database platform or environment, this text fully reflects the field's latest realities and challenges. Drawing on more than thirty years of database experience, Mullins focuses on problems that today's DBAs actually face, and skills and knowledge they simply must have. Mullins presents realistic, thorough, and up-to-date coverage of every DBA task, including creating database environments, data modeling, normalization, design, performance, data integrity, compliance, governance, security, backup/recovery, disaster planning, data and storage management, data movement/distribution, data warehousing, connectivity, metadata, tools, and more. This edition adds new coverage of "Big Data," database appliances, cloud computing, and NoSQL. Mullins includes an entirely new chapter on the DBA's role in regulatory compliance, with substantial new material on data breaches, auditing, encryption, retention, and metadata management. You'll also find an all-new glossary, plus up-to-the-minute DBA rules of thumb.

In this book, cofounder and lead developer James Gardner brings you a comprehensive introduction to Pylons, the web framework that uses the best of Ruby, Python, and Perl and the emerging WSGI standard to provide structure and flexibility. You'll learn how to create your own Pylons-driven web site and attain the mastery of advanced Pylons features. You'll also learn how to stretch Pylons to its fullest ability, as well as share Gardner's unique insight and extensive experience in developing and deploying Pylons for a wide variety of situations.

This second edition of Michael Kofler's acclaimed MySQL book has updated and expanded to cover MySQL 4.0, the most recent production release of the popular open source database, which boasts more than 4 million users worldwide. Like the first edition, this revision, which has been renamed to reflect the breadth and depth of Kofler's coverage of the topic, provides a thorough introduction to the installation, configuration, implementation, and administration of MySQL. In addition, Kofler demonstrates how you can use MySQL in conjunction with various other technologies to create database-driven websites, and he gives practical advice on database design. Kofler also covers what's coming up next in MySQL 4.1.

Summary MongoDB in Action, Second Edition is a completely revised and updated version. It introduces MongoDB 3.0 and the document-oriented database model. This perfectly paced book gives you both the big picture you'll need as a developer and enough low-level detail to satisfy system engineers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology This document-oriented database was built for high availability, supports rich, dynamic schemas, and lets you easily distribute data across multiple servers. MongoDB 3.0 is flexible, scalable, and very fast, even with big data loads. About the Book MongoDB in Action, Second Edition is a completely revised and updated version. It introduces MongoDB 3.0 and the document-oriented database model. This perfectly paced book gives you both the big picture you'll need as a developer and enough low-level detail to satisfy system engineers. Lots of examples will help you develop confidence in the crucial area of data modeling. You'll also love the deep explanations of each feature, including replication, auto-sharding, and deployment. What's Inside Indexes, queries, and standard DB operations Aggregation and text searching Map-reduce for custom aggregations and reporting Deploying for scale and high availability Updated for Mongo 3.0 About the Reader Written for developers. No previous MongoDB or NoSQL experience is assumed. About the Authors After working at MongoDB, Kyle Banker is now at a startup. Peter Bakkum is a developer with MongoDB expertise. Shaun Verch has worked on the core server team at MongoDB. A Genentech engineer, Doug Garrett is one of the winners of the MongoDB Innovation Award for Analytics. A software architect, Tim Hawkins has led search engineering at Yahoo Europe. Technical Contributor: Wouter Thielen. Technical Editor: Mihalis Tsoukalos. Table of Contents PART 1 GETTING STARTED A database for the modern web MongoDB through the JavaScript shell Writing programs using MongoDB PART 2 APPLICATION DEVELOPMENT IN MONGODB Document-oriented data Constructing queries Aggregation Updates, atomic operations, and deletes PART 3 MONGODB MASTERY Indexing and query optimization Text search WiredTiger and pluggable storage Replication Scaling your system with sharding Deployment and administration

An introduction to the engineering principles of embedded systems, with a focus on modeling, design, and analysis of cyber-

physical systems. The most visible use of computers and software is processing information for human consumption. The vast majority of computers in use, however, are much less visible. They run the engine, brakes, seatbelts, airbag, and audio system in your car. They digitally encode your voice and construct a radio signal to send it from your cell phone to a base station. They command robots on a factory floor, power generation in a power plant, processes in a chemical plant, and traffic lights in a city. These less visible computers are called embedded systems, and the software they run is called embedded software. The principal challenges in designing and analyzing embedded systems stem from their interaction with physical processes. This book takes a cyber-physical approach to embedded systems, introducing the engineering concepts underlying embedded systems as a technology and as a subject of study. The focus is on modeling, design, and analysis of cyber-physical systems, which integrate computation, networking, and physical processes. The second edition offers two new chapters, several new exercises, and other improvements. The book can be used as a textbook at the advanced undergraduate or introductory graduate level and as a professional reference for practicing engineers and computer scientists. Readers should have some familiarity with machine structures, computer programming, basic discrete mathematics and algorithms, and signals and systems.

Advanced Topics

The Java EE 7 Tutorial

Data Structures and Algorithms in Java

Concepts and Implementation

Joe Celko's Complete Guide to NoSQL

Started by small group of well known scientists with the aim of sharing knowledge, experiences, and results on all aspects of cluster computing, the initiative of a workshop on cluster computing received more attention after IFIP WG 10.3 and IEEE Romania Section accepted our request for sponsorship. Moreover, the application for a NATO ARW grant was successful, leading to a greater interest in the workshop. In this respect, we have to say that we chose Romania in order to attract scientists from Central and Eastern European countries and improve the cooperation in the region, in the field of cluster computing. We had an extremely short time to organize the event, but many people joined us and enthusiastically contributed to the process. The success of the workshop is wholly due to the hard work of the organizing committee, members of the program committee, key speakers, speakers from industry, and authors of accepted papers. The workshop consisted of invited and regular paper presentations, followed by discussions, on many important current and emerging topics ranging from scheduling and load balancing to grids. The key speakers devoted their time and efforts to presenting the most interesting results of their research groups, and we all thank them for this . All papers were peer reviewed by two or three reviewers.

The Definitive Guide to GCC

Readings in Database Systems

NATO Advanced Research Workshop, IWCC 2001, Mangalia, Romania, September 1-6, 2001. Revised Papers

The Berkeley DB Book

Oracle NoSQL Database