

Read Book Tutorial Berkeley Db Java Edition Pdf For Free

The Definitive Guide to Berkeley DB XML Berkeley DB The Berkeley DB Book Oracle NoSQL Database The CIO's Guide to Oracle Products and Solutions Oracle Essentials InfoWorld JAVA Developer's Journal Software Composition Building Search Applications Databases on Linux Advances in Databases and Information Systems Pattern and Data Analysis in Healthcare Settings Advances in Databases Processing XML Documents with Oracle JDeveloper 11g Database Systems for Advanced Applications On the Move to Meaningful Internet Systems: OTM 2009 Databases, Information Systems, and Peer-to-Peer Computing Advances in Swarm Intelligence Dr. Dobb's Journal NoSQL For Dummies Scalable Uncertainty Management Mastering DynamoDB The Definitive Guide to Berkeley DB XML Java Persistence for Relational Databases Data Structures and Algorithms in Java XML Data Management New Directions in Web Data Management 1 Pro EJB 3 Dr. Dobb's Journal of Software Tools for the Professional Programmer Beginning Java 7 The Definitive Guide to db4o Understanding MySQL Internals Lucene in Action Feature-Oriented Software Product Lines Pro Hibernate and MongoDB InfoWorld Spark: The Definitive Guide Professional NoSQL The Definitive Guide to SQLite

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. This book constitutes the refereed proceedings of the 9th East European Conference on Advances in Databases and Information Systems, ADBIS 2005, held in Tallinn, Estonia, in September 2005. The 27 revised full papers presented together with an invited paper were carefully reviewed and selected from 144 submissions. The papers are organized in topical sections on database theory, database modelling and physical database design, query processing, heterogeneous databases and interoperability, XML and databases, data mining and knowledge discovery, information systems and software engineering, and information systems development. Beginning Java 7 guides you through version 7 of the Java language and a wide assortment of platform APIs. New Java 7 language features that are discussed include switch-on-string and try-with-resources. APIs that are discussed include Threading, the Collections Framework, the Concurrency Utilities, Swing, Java 2D, networking, JDBC, SAX, DOM, StAX, XPath, JAX-WS, and SAAJ. This book also presents an introduction to Android app development so that you can apply some of its knowledge to the exciting world of Android app development. This book presents the

following table of contents: Chapter 1 introduces you to Java and begins to cover the Java language by focusing on fundamental concepts such as comments, identifiers, variables, expressions, and statements. Chapter 2 continues to explore this language by presenting all of its features for working with classes and objects. You learn about features related to class declaration and object creation, encapsulation, information hiding, inheritance, polymorphism, interfaces, and garbage collection. Chapter 3 focuses on the more advanced language features related to nested classes, packages, static imports, exceptions, assertions, annotations, generics, and enums. Additional chapters introduce you to the few features not covered in Chapters 1 through 3. Chapter 4 largely moves away from covering language features (although it does introduce class literals and strictfp) while focusing on language-oriented APIs. You learn about Math, StrictMath, Package, Primitive Type Wrapper Classes, Reference, Reflection, String, StringBuffer and StringBuilder, Threading, BigDecimal, and BigInteger in this chapter. Chapter 5 begins to explore Java's utility APIs by focusing largely on the Collections Framework. However, it also discusses legacy collection-oriented APIs and how to create your own collections. Chapter 6 continues to focus on utility APIs by presenting the concurrency utilities along with the Objects and Random classes. Chapter 7 moves you away from the command-line user interfaces that appear in previous chapters and toward graphical user interfaces. You first learn about the Abstract Window Toolkit foundation, and then explore the Java Foundation Classes in terms of Swing and Java 2D. Appendix C explores Accessibility and Drag and Drop. Chapter 8 explores filesystem-oriented I/O in terms of the File, RandomAccessFile, stream, and writer/reader classes. Chapter 9 introduces you to Java's network APIs (e.g., sockets). It also introduces you to the JDBC API for interacting with databases along with the Java DB database product. Chapter 10 dives into Java's XML support by first presenting an introduction to XML (including DTDs and schemas). It next explores the SAX, DOM, StAX, XPath, and XSLT APIs. It even briefly touches on the Validation API. While exploring XPath, you encounter namespace contexts, extension functions and function resolvers, and variables and variable resolvers. Chapter 11 introduces you to Java's support for SOAP-based and RESTful web services. As well as providing you with the basics of these web service categories, Chapter 11 presents some advanced topics, such as working with the SAAJ API to communicate with a SOAP-based web service without having to rely on JAX-WS. You will appreciate having learned about XML in Chapter 10 before diving into this chapter. Chapter 12 helps you put to use some of the knowledge you've gathered in previous chapters by showing you how to use Java to write an Android app's source code. This chapter introduces you to Android, discusses its architecture, shows you how to install necessary tools, and develops a simple app. Appendix A presents the solutions to the programming exercises that appear near the end of Chapters 1 through 12. Appendix B introduces you to Java's Scripting API along

with Java 7's support for dynamically typed languages. Appendix C introduces you to additional APIs and architecture topics. Examples include Accessibility, classloaders, Console, Drag and Drop, Java Native Interface, and System Tray. Appendix D presents a gallery of significant applications that demonstrate various aspects of Java. Unfortunately, there are limits to how much knowledge can be crammed into a print book. For this reason, Appendixes A, B, C, and D are not included in this book's pages. Instead, these appendixes are freely distributed as PDF files. Appendixes A and B are bundled with the book's associated code file at the Apress website (<http://www.apress.com/9781430239093>). Appendixes C and D are bundled with their respective code files at my TutorTutor.ca website (<http://tutortutor.ca/cgi-bin/makepage.cgi?/books/bj7>). This book constitutes the workshop proceedings of the 15th International Conference on Database Systems for Advanced Applications, DASFAA 2010, held in Tsukuba, Japan, in April 2010. The volume contains six workshops, each focusing on specific research issues that contribute to the main themes of the DASFAA conference: The First International Workshop on Graph Data Management: Techniques and Applications (GDM 2010), The Second International Workshop on Benchmarking of Database Management Systems and Data-Oriented Web Technologies (BenchmarkX'10); The Third International Workshop on Managing Data Quality in Collaborative Information Systems (MCIS2010), The Workshop on Social Networks and Social Media Mining on the Web (SNSMW2010), The Data Intensive eScience Workshop (DIEW 2010), and The Second International Workshop on Ubiquitous Data Management (UDM2010). 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. InfoWorld is targeted to Senior IT professionals. Content is segmented into Channels and Topic Centers. InfoWorld also celebrates people, companies, and projects. 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! Business and medical professionals rely on large data sets to identify trends or other knowledge that can be gleaned from the collection of it. New technologies concentrate on data’s management, but do not facilitate users’ extraction of meaningful outcomes. Pattern and Data Analysis in Healthcare Settings investigates the approaches to shift computing from analysis on-demand to knowledge on-demand. By providing innovative tactics to apply data and pattern analysis, these practices are optimized into pragmatic sources of knowledge for healthcare professionals. This publication is an exhaustive source for policy makers, developers, business professionals, healthcare providers, and graduate students concerned with data retrieval and analysis. First EJB 3.0 book on the market and a definitive guide to the major innovation in EJB: the new persistence API Offers unparalleled insight and expertise: lead authored by the co-lead on the EJB 3.0 spec (Mike Keith) 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. Sperko focuses on the overall problem of how to store the primary

component of any Java application, the Java object, in the most common business tool: the relational database. Learn how to use, deploy, and maintain Apache Spark with this comprehensive guide, written by the creators of the open-source cluster-computing framework. With an emphasis on improvements and new features in Spark 2.0, authors Bill Chambers and Matei Zaharia break down Spark topics into distinct sections, each with unique goals. You'll explore the basic operations and common functions of Spark's structured APIs, as well as Structured Streaming, a new high-level API for building end-to-end streaming applications. Developers and system administrators will learn the fundamentals of monitoring, tuning, and debugging Spark, and explore machine learning techniques and scenarios for employing MLlib, Spark's scalable machine-learning library. Get a gentle overview of big data and Spark Learn about DataFrames, SQL, and Datasets—Spark's core APIs—through worked examples Dive into Spark's low-level APIs, RDDs, and execution of SQL and DataFrames Understand how Spark runs on a cluster Debug, monitor, and tune Spark clusters and applications Learn the power of Structured Streaming, Spark's stream-processing engine Learn how you can apply MLlib to a variety of problems, including classification or recommendation The Berkeley DB Book is intended to be a practical guide to the intricacies of Berkeley DB; an in-depth analysis of the complex design issues which are often covered in terse footnotes in the dense Berkeley DB reference manual. It explains the technology at a higher level and also covers the internals with generous code and design examples. Berkeley DB is becoming the database of choice for appliance makers and for in memory cache of large scale applications like search engines and high traffic web sites. This book constitutes the refereed proceedings of the 8th International Symposium on Software Composition, SC 2009, held in Zurich, Switzerland, in July 2009. The workshop has been organized as an event co-located with the TOOLS Europe 2009 conference. The 10 revised full papers presented together with 2 invited lectures were carefully reviewed and selected from 34 submissions. The papers reflect current research in software composition to foster developing of composition models and techniques by using aspect- and service-oriented programming, specification of component contracts and protocols, methods of correct components composition, as well as verification, validation and testing techniques - even in pervasive computing environments and for the Web. 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. In this book, you will find discussions on the newest native XML databases, along with information on working with XML-enabled relational database systems. In addition, XML Data Management thoroughly examines benchmarks and analysis techniques for performance of XML databases. This book is best used by students that are knowledgeable in database technology and are familiar with XML. 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. Internet-based information systems, the second covering the large-scale integration of heterogeneous computing systems and data resources with the aim of providing a global computing space.

Each of these four conferences encourages researchers to treat their respective topics within a framework that incorporates jointly (a) theory, (b) conceptual design and development, and (c) applications, in particular case studies and industrial solutions. Following and expanding the model created in 2003, we again solicited and selected quality workshop proposals to complement the more "archival" nature of the main conferences with research results in a number of selected and more "avant-garde" areas related to the general topic of Web-based distributed computing. For instance, the so-called Semantic Web has given rise to several novel research areas combining linguistics, information systems technology, and artificial intelligence, such as the modeling of (legal) regulatory systems and the ubiquitous nature of their usage. We were glad to see that ten of our earlier successful workshops (ADI, CAMS, EI2N, SWWS, ORM, OnToContent, MONET, SEMELS, COMBEK, IWSSA) re-appeared in 2008 with a second, third or even fourth edition, sometimes by alliance with other newly emerging workshops, and that no fewer than three brand-new independent workshops could be selected from proposals and hosted: ISDE, ODIS and Beyond SAWSDL. Workshop attendees productively mingled with each other and with those of the main conferences, and there was considerable overlap in authors. Get up to speed on the nuances of NoSQL databases and what they mean for your organization This easy to read guide to NoSQL databases provides the type of no-nonsense overview and analysis that you need to learn, including what NoSQL is and which database is right for you. Featuring specific evaluation criteria for NoSQL databases, along with a look into the pros and cons of the most popular options, NoSQL For Dummies provides the fastest and easiest way to dive into the details of this incredible technology. You'll gain an understanding of how to use NoSQL databases for mission-critical enterprise architectures and projects, and real-world examples reinforce the primary points to create an action-oriented resource for IT pros. If you're planning a big data project or platform, you probably already know you need to select a NoSQL database to complete your architecture. But with options flooding the market and updates and add-ons coming at a rapid pace, determining what you require now, and in the future, can be a tall task. This is where NoSQL For Dummies comes in!

Learn the basic tenets of NoSQL databases and why they have come to the forefront as data has outpaced the capabilities of relational databases Discover major players among NoSQL databases, including Cassandra, MongoDB, MarkLogic, Neo4J, and others Get an in-depth look at the benefits and disadvantages of the wide variety of NoSQL database options Explore the needs of your organization as they relate to the capabilities of specific NoSQL databases Big data and Hadoop get all the attention, but when it comes down to it, NoSQL databases are the engines that power many big data analytics initiatives. With NoSQL For Dummies, you'll go beyond relational databases to ramp up your enterprise's data architecture in no time. This book constitutes the thoroughly refereed post-conference proceedings of the 28th British National Conference on Databases, BNCOD 28, held in Manchester, UK, in July 2011. The 13 revised full papers, 2 short papers, 2 demo papers and 1 poster paper presented together with the abstracts of 2 keynote talks and 1 tutorial paper were carefully reviewed and selected from 44 submissions. The papers cover a wide range of topics such as XML compression, XML updates, column-oriented stores, provenance, warehousing, streamed data, data mashups, dataspace, sensor network query processing, and pattern-oriented search. 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. Hibernate and MongoDB are a powerful combination of open source persistence and NoSQL technologies for today's Java-based enterprise and cloud application developers. Hibernate is the leading open source Java-based persistence, object relational management engine, recently repositioned as an object grid management engine. MongoDB is a growing, popular open source NoSQL framework, especially popular among cloud application and big data developers. With these two, enterprise and cloud developers have a "complete out of the box" solution. Pro Hibernate and MongoDB shows you how to use and integrate Hibernate and MongoDB. More specifically, this book guides you through the bootstrap; building transactions; handling queries and query entities; and mappings. Then, this book explores the principles and techniques for taking these application principles to the cloud, using the OpenShift Platform as a Service (PaaS) and more. In this book, you get two case studies: An enterprise application using Hibernate and MongoDB. then, A cloud application (OpenShip)

migrated from the enterprise application case study After reading or using this book, you come away with the experience from two case studies that give you possible frameworks or templates that you can apply to your own specific application or cloud application building context. 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. 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. Oracle is an enormous system, with myriad technologies, options, and releases. Most users--even experienced developers and database administrators--find it difficult to get a handle on the full scope of the Oracle database. And, as each new Oracle version is released, users find themselves under increasing pressure to learn about a whole range of new technologies. The latest challenge is Oracle Database 11g. This book distills an enormous amount of information about Oracle into a compact, easy-to-read volume filled with focused text, illustrations, and helpful hints. It contains chapters on: Oracle products, options, data structures, and overall architecture for Oracle Database 11g, as well as earlier releases (Oracle Database 10g, Oracle9i, and Oracle8i) Installing, running, managing, monitoring, networking, and tuning Oracle, including Enterprise Manager (EM) and Oracle's self-tuning and management capabilities; and using Oracle security, auditing, and compliance (a new chapter in this edition) Multiuser concurrency, data warehouses, distributed databases, online transaction processing (OLTP), high availability, and hardware architectures (e.g.,

SMP, clusters, NUMA, and grid computing) Features beyond the Oracle database: Oracle Application Express, Fusion Middleware (including Oracle Application Server), and database SOA support as a Web services provider The latest Oracle Database 11g features: query result set caching, Automatic Memory Management, the Real Application Testing, Advanced Compression, Total Recall, and Active Data Guard Option Options, changes to the OLAP Option (transparently accessed and managed as materialized views), the Flashback transaction command, transparent data encryption, the Support Workbench (and diagnosability infrastructure), and partitioning enhancements (including interval and new composite types) For new Oracle users, DBAs, developers, and managers, Oracle Essentials provides an invaluable, all-in-one introduction to the full range of Oracle features and technologies, including the just-released Oracle Database 11g features. But even if you already have a library full of Oracle documentation, you'll find that this compact book is the one you turn to, again and again, as your one-stop, truly essential reference. "Oracle Essentials gives a clear explanation of the key database concepts and architecture underlying the Oracle database. It's a great reference for anyone doing development or management of Oracle databases." --Andrew Mendelsohn, Senior Vice President, Database Server Technologies, Oracle Corporation

Lucene is a gem in the open-source world Lucene in Action'" is the authoritative guide to Lucene. It describes how to index your data, including types you definitely need to know such as MS Word, PDF, HTML, and XML. It introduces you to searching, sorting, filtering, and highlighting search results. Lucene powers search in surprising places What's Inside - How to integrate Lucene into your applications - Ready-to-use framework for rich document handling - Case studies including Nutch, TheServerSide, jGuru, etc. - Lucene ports to Perl, Python, C#/.Net, and C++ - Sorting, filtering, term vectors, multiple, and remote index searching - The new SpanQuery family, extending query parser, hit collecting - Performance testing and tuning - Lucene add-ons (hit highlighting, synonym lookup, and others) 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. Lucene, LingPipe, and Gate are popular open source tools to build powerful search applications. Building Search Applications describes functions from Lucene that include indexing, searching, ranking, and spelling correction to build search engines. With this book you will learn to: Extract tokens from text using custom tokenizers and analyzers from Lucene, LingPipe, and Gate. Construct

a search engine index with an optional backend database to manage large document collections. Explore the wide range of Lucene queries to search an index, understand the ranking algorithm for a query, and suggest spelling corrections. Find the names of people, places, and other entities in text using LingPipe and Gate. Categorize documents by topic using classifiers and build groups of self-organized documents using clustering algorithms from LingPipe. Create a Web crawler to scan the Web, Intranet, or desktop using Nutch. Track the sentiment of articles published on the Web with LingPipe. Small, special-purpose computing devices and high-end core Internet servers need fast, reliable database management. Berkeley DB is an embedded database that provides high-performance, scalable, transaction-protected and recoverable data management services to applications. Extremely portable, this library runs under almost all UNIX and Windows variants, as well as a number of embedded, real-time operating systems. Berkeley DB is the ultimate resource for the world's most widely deployed embedded database engine. This book will aid software architects and engineers, product managers, and systems and network administrators without the overhead imposed by other database products. Designed by programmers for programmers, this classic library style toolkit provides a broad base of functionality to application writers. This book will help you to make intelligent choices about when and how to use Berkeley DB to meet your needs. You can visit the Sleepycat website to get the latest errata for this book. NOTE: The first printing of this book contained an error in the table of contents that caused the page numbers to be off. This will be corrected in the second printing. If you have an earlier edition, you can download a pdf of the correct table of contents that you can print out and use with your book. If you have any questions, please feel free to contact the editor of this book at stephanie.wall@newriders.com. This book and its companion volume, LNCS vols. 7331 and 7332, constitute the Proceedings of the Third International conference on Swarm Intelligence, ICSI 2012, held in Shenzhen, China in June 2012. The 145 full papers presented were carefully reviewed and selected from 247 submissions. The papers are organized in 27 cohesive sections covering all major topics of swarm intelligence research and developments. Managing uncertainty and inconsistency has been extensively explored in - ticial Intelligence over a number of years. Now with the advent of massive amounts of data and knowledge from distributed heterogeneous, and potentially con?icting, sources, there is interest in developing and applying formalisms for uncertainty and inconsistency widely in systems that need to better manage this data and knowledge. The annual International Conference on Scalable Uncertainty Management (SUM) has grown out of this wide-ranging interest in managing uncertainty and inconsistency in databases, the Web, the Semantic Web, and AI. It aims at bringing together all those interested in the management of large volumes of uncertainty and inconsistency, irrespective of whether they are in databases, the Web, the Semantic Web, or in AI, as well as in other areas such as information retrieval, risk

analysis, and computer vision, where significant computational efforts are needed. After a promising First International Conference on Scalable Uncertainty Management was held in Washington DC, USA in 2007, the conference series has been successfully held in Napoli, Italy, in 2008, and again in Washington DC, USA, in 2009. Create, validate, and transform XML documents with Oracle's JDeveloper IDE using this book and eBook. Why store data in text files? It is better to use a database. More and more often data is supported by programs and stored in databases. Depending on needs these can be simple file bases or complex relational databases. The micro-course gives basic information about databases used in the Linux system. It discusses Berkeley DB and SQLite file bases. Keywords: database, Berkeley db, SQLite Looks at the workings of MySQL 5, covering such topics as configuration variables, storage engines, the table lock manager, the communication protocol, and server classes. This book addresses the major issues in the Web data management related to technologies and infrastructures, methodologies and techniques as well as applications and implementations. Emphasis is placed on Web engineering and technologies, Web graph managing, searching and querying and the importance of social Web. If you have interest in DynamoDB and want to know what DynamoDB is all about and become proficient in using it, this is the book for you. If you are an intermediate user who wishes to enhance your knowledge of DynamoDB, this book is aimed at you. Basic familiarity with programming, NoSQL, and cloud computing concepts would be helpful. InfoWorld is targeted to Senior IT professionals. Content is segmented into Channels and Topic Centers. InfoWorld also celebrates people, companies, and projects.

digitaltutorials.jrn.columbia.edu