

---

# Access Free Pdf Servlets Core Introduction Hadoop

---

Recognizing the quirk ways to acquire this ebook **Pdf Servlets Core Introduction Hadoop** is additionally useful. You have remained in right site to start getting this info. acquire the Pdf Servlets Core Introduction Hadoop associate that we come up with the money for here and check out the link.

You could purchase lead Pdf Servlets Core Introduction Hadoop or get it as soon as feasible. You could speedily download this Pdf Servlets Core Introduction Hadoop after getting deal. So, with you require the ebook swiftly, you can straight get it. Its therefore entirely easy and consequently fats, isnt it? You have to favor to in this flavor

---

## KEY=CORE - MATTEO MELINA

---

---

### APACHE SPARK FOR THE ENTERPRISE: SETTING THE BUSINESS FREE

---

IBM Redbooks **Analytics is increasingly an integral part of day-to-day operations at today's leading businesses, and transformation is also occurring through huge growth in mobile and digital channels. Enterprise organizations are attempting to leverage analytics in new ways and transition existing analytics capabilities to respond with more flexibility while making the most efficient use of highly valuable data science skills. The recent growth and adoption of Apache Spark as an analytics framework and platform is very timely and helps meet these challenging demands. The Apache Spark environment on IBM z/OS® and Linux on IBM z Systems™ platforms allows this analytics framework to run on the same enterprise platform as the originating sources of data and transactions that feed it. If most of the data that will be used for Apache Spark analytics, or the most sensitive or quickly changing data is originating on z/OS, then an Apache Spark z/OS based environment will be the optimal choice for performance, security, and governance. This IBM® Redpaper™ publication explores the enterprise analytics market, use of Apache Spark on IBM z Systems™ platforms, integration between Apache Spark and other enterprise data sources, and case studies and examples of what can be achieved with Apache Spark in enterprise environments. It is of interest to data scientists, data engineers, enterprise architects, or anybody looking to better understand how to combine an analytics framework and platform on enterprise systems.**

---

## HBASE

---

---

### THE DEFINITIVE GUIDE

---

"O'Reilly Media, Inc." **If your organization is looking for a storage solution to accommodate a virtually endless amount of data, this book will show you how Apache HBase can fulfill your needs. As the open source implementation of Google's BigTable architecture, HBase scales to billions of rows and millions of columns, while ensuring that write and read performance remain constant. HBase: The Definitive Guide provides the details you require, whether you simply want to evaluate this high-performance, non-relational database, or put it into practice right away. HBase's adoption rate is beginning to climb, and several IT executives are asking pointed questions about this high-capacity database. This is the only book available to give you meaningful answers. Learn how to distribute large datasets across an inexpensive cluster of commodity servers Develop HBase clients in many programming languages, including Java, Python, and Ruby Get details on HBase's primary storage system, HDFS—Hadoop's distributed and replicated filesystem Learn how HBase's native interface to Hadoop's MapReduce framework enables easy development and execution of batch jobs that can scan entire tables Discover the integration between HBase and other facets of the Apache Hadoop project**

---

## PROFESSIONAL NOSQL

---

John Wiley & Sons **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.**

---

## HADOOP: THE DEFINITIVE GUIDE

---

"O'Reilly Media, Inc." **Ready to unlock the power of your data? With this comprehensive guide, you'll learn how to build and maintain reliable, scalable, distributed systems with Apache Hadoop. This book is ideal for programmers looking to analyze datasets of any size, and for administrators who want to set up and run Hadoop clusters. You'll find illuminating case studies that demonstrate how Hadoop is used to solve specific problems. This third edition covers**

recent changes to Hadoop, including material on the new MapReduce API, as well as MapReduce 2 and its more flexible execution model (YARN). Store large datasets with the Hadoop Distributed File System (HDFS) Run distributed computations with MapReduce Use Hadoop's data and I/O building blocks for compression, data integrity, serialization (including Avro), and persistence Discover common pitfalls and advanced features for writing real-world MapReduce programs Design, build, and administer a dedicated Hadoop cluster—or run Hadoop in the cloud Load data from relational databases into HDFS, using Sqoop Perform large-scale data processing with the Pig query language Analyze datasets with Hive, Hadoop's data warehousing system Take advantage of HBase for structured and semi-structured data, and ZooKeeper for building distributed systems

---

## HADOOP IN ACTION

---

[Simon and Schuster](#) **Hadoop in Action** teaches readers how to use Hadoop and write MapReduce programs. The intended readers are programmers, architects, and project managers who have to process large amounts of data offline. **Hadoop in Action** will lead the reader from obtaining a copy of Hadoop to setting it up in a cluster and writing data analytic programs. The book begins by making the basic idea of Hadoop and MapReduce easier to grasp by applying the default Hadoop installation to a few easy-to-follow tasks, such as analyzing changes in word frequency across a body of documents. The book continues through the basic concepts of MapReduce applications developed using Hadoop, including a close look at framework components, use of Hadoop for a variety of data analysis tasks, and numerous examples of Hadoop in action. **Hadoop in Action** will explain how to use Hadoop and present design patterns and practices of programming MapReduce. MapReduce is a complex idea both conceptually and in its implementation, and Hadoop users are challenged to learn all the knobs and levers for running Hadoop. This book takes you beyond the mechanics of running Hadoop, teaching you to write meaningful programs in a MapReduce framework. This book assumes the reader will have a basic familiarity with Java, as most code examples will be written in Java. Familiarity with basic statistical concepts (e.g. histogram, correlation) will help the reader appreciate the more advanced data processing examples. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book.

---

## SERVLET & JSP: A TUTORIAL, SECOND EDITION

---

[Brainy Software Inc](#) **Servlet and JavaServer Pages (JSP)** are the underlying technologies for developing web applications in Java. They are essential for any programmer to master in order to effectively use frameworks such as JavaServer Faces, Struts 2 or Spring MVC. Covering Servlet 3.1 and JSP 2.3, this book explains the important programming concepts and design models in Java web development as well as related technologies and new features in the latest versions of Servlet and JSP. With comprehensive coverage and a lot of examples, this book is a guide to building real-world applications.

---

## DISTRIBUTED TRACING IN PRACTICE

---



---

## INSTRUMENTING, ANALYZING, AND DEBUGGING MICROSERVICES

---

[O'Reilly Media](#) Most applications today are distributed in some fashion. Monitoring the health and performance of these distributed architectures requires a new approach. Enter distributed tracing, a method of profiling and monitoring applications—especially those that use microservice architectures. There's just one problem: distributed tracing can be hard. But it doesn't have to be. With this practical guide, you'll learn what distributed tracing is and how to use it to understand the performance and operation of your software. Key players at Lightstep walk you through instrumenting your code for tracing, collecting the data that your instrumentation produces, and turning it into useful, operational insights. If you want to start implementing distributed tracing, this book tells you what you need to know. You'll learn: The pieces of a distributed tracing deployment: Instrumentation, data collection, and delivering value Best practices for instrumentation (the methods for generating trace data from your service) How to deal with or avoid overhead, costs, and sampling How to work with spans (the building blocks of request-based distributed traces) and choose span characteristics that lead to valuable traces Where distributed tracing is headed in the future

---

## SOLR IN ACTION

---

[Simon and Schuster](#) **Summary Solr in Action** is a comprehensive guide to implementing scalable search using Apache Solr. This clearly written book walks you through well-documented examples ranging from basic keyword searching to scaling a system for billions of documents and queries. It will give you a deep understanding of how to implement core Solr capabilities. **About the Book** Whether you're handling big (or small) data, managing documents, or building a website, it is important to be able to quickly search through your content and discover meaning in it. Apache Solr is your tool: a ready-to-deploy, Lucene-based, open source, full-text search engine. Solr can scale across many servers to enable real-time queries and data analytics across billions of documents. **Solr in Action** teaches you to implement scalable search using Apache Solr. This easy-to-read guide balances conceptual discussions with practical examples to show you how to implement all of Solr's core capabilities. You'll master topics like text analysis, faceted search, hit highlighting, result grouping, query suggestions, multilingual search, advanced geospatial and data operations, and relevancy tuning. This book assumes basic knowledge of Java and standard database technology. No prior knowledge of Solr or Lucene is required. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. **What's Inside** How to scale Solr for big data Rich real-world examples Solr as a NoSQL data store Advanced multilingual, data, and relevancy tricks Coverage of versions through Solr 4.7 **About the Authors** Trey

Grainger is a director of engineering at CareerBuilder. Timothy Potter is a senior member of the engineering team at LucidWorks. The authors work on the scalability and reliability of Solr, as well as on recommendation engine and big data analytics technologies. Table of Contents PART 1 MEET SOLR Introduction to Solr Getting to know Solr Key Solr concepts Configuring Solr Indexing Text analysis PART 2 CORE SOLR CAPABILITIES Performing queries and handling results Faceted search Hit highlighting Query suggestions Result grouping/field collapsing Taking Solr to production PART 3 TAKING SOLR TO THE NEXT LEVEL SolrCloud Multilingual search Complex query operations Mastering relevancy

---

## MACHINE LEARNING

---

### HANDS-ON FOR DEVELOPERS AND TECHNICAL PROFESSIONALS

---

[John Wiley & Sons](#) Dig deep into the data with a hands-on guide to machine learning with updated examples and more! **Machine Learning: Hands-On for Developers and Technical Professionals** provides hands-on instruction and fully-coded working examples for the most common machine learning techniques used by developers and technical professionals. The book contains a breakdown of each ML variant, explaining how it works and how it is used within certain industries, allowing readers to incorporate the presented techniques into their own work as they follow along. A core tenant of machine learning is a strong focus on data preparation, and a full exploration of the various types of learning algorithms illustrates how the proper tools can help any developer extract information and insights from existing data. The book includes a full complement of Instructor's Materials to facilitate use in the classroom, making this resource useful for students and as a professional reference. At its core, machine learning is a mathematical, algorithm-based technology that forms the basis of historical data mining and modern big data science. Scientific analysis of big data requires a working knowledge of machine learning, which forms predictions based on known properties learned from training data. Machine Learning is an accessible, comprehensive guide for the non-mathematician, providing clear guidance that allows readers to: Learn the languages of machine learning including Hadoop, Mahout, and Weka Understand decision trees, Bayesian networks, and artificial neural networks Implement Association Rule, Real Time, and Batch learning Develop a strategic plan for safe, effective, and efficient machine learning By learning to construct a system that can learn from data, readers can increase their utility across industries. Machine learning sits at the core of deep dive data analysis and visualization, which is increasingly in demand as companies discover the goldmine hiding in their existing data. For the tech professional involved in data science, Machine Learning: Hands-On for Developers and Technical Professionals provides the skills and techniques required to dig deeper.

---

## APACHE OOZIE

---

### THE WORKFLOW SCHEDULER FOR HADOOP

---

["O'Reilly Media, Inc."](#) Get a solid grounding in Apache Oozie, the workflow scheduler system for managing Hadoop jobs. With this hands-on guide, two experienced Hadoop practitioners walk you through the intricacies of this powerful and flexible platform, with numerous examples and real-world use cases. Once you set up your Oozie server, you'll dive into techniques for writing and coordinating workflows, and learn how to write complex data pipelines. Advanced topics show you how to handle shared libraries in Oozie, as well as how to implement and manage Oozie's security capabilities. Install and configure an Oozie server, and get an overview of basic concepts Journey through the world of writing and configuring workflows Learn how the Oozie coordinator schedules and executes workflows based on triggers Understand how Oozie manages data dependencies Use Oozie bundles to package several coordinator apps into a data pipeline Learn about security features and shared library management Implement custom extensions and write your own EL functions and actions Debug workflows and manage Oozie's operational details

---

## HADOOP IN PRACTICE

---

[Simon and Schuster](#) Summary **Hadoop in Practice, Second Edition** provides over 100 tested, instantly useful techniques that will help you conquer big data, using Hadoop. This revised new edition covers changes and new features in the Hadoop core architecture, including MapReduce 2. Brand new chapters cover YARN and integrating Kafka, Impala, and Spark SQL with Hadoop. You'll also get new and updated techniques for Flume, Sqoop, and Mahout, all of which have seen major new versions recently. In short, this is the most practical, up-to-date coverage of Hadoop available anywhere. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Book It's always a good time to upgrade your Hadoop skills! **Hadoop in Practice, Second Edition** provides a collection of 104 tested, instantly useful techniques for analyzing real-time streams, moving data securely, machine learning, managing large-scale clusters, and taming big data using Hadoop. This completely revised edition covers changes and new features in Hadoop core, including MapReduce 2 and YARN. You'll pick up hands-on best practices for integrating Spark, Kafka, and Impala with Hadoop, and get new and updated techniques for the latest versions of Flume, Sqoop, and Mahout. In short, this is the most practical, up-to-date coverage of Hadoop available. Readers need to know a programming language like Java and have basic familiarity with Hadoop. What's Inside Thoroughly updated for Hadoop 2 How to write YARN applications Integrate real-time technologies like Storm, Impala, and Spark Predictive analytics using Mahout and RR Readers need to know a programming language like Java and have basic familiarity with Hadoop. About the Author Alex Holmes works on tough big-data problems. He is a software engineer, author, speaker, and blogger specializing in large-scale Hadoop projects. Table of Contents PART 1 BACKGROUND AND FUNDAMENTALS Hadoop in a heartbeat Introduction to YARN PART 2 DATA LOGISTICS Data serialization—working with text and beyond Organizing and optimizing data in HDFS Moving data into and out of Hadoop PART 3 BIG DATA PATTERNS Applying MapReduce patterns to big data Utilizing data structures and algorithms

at scale Tuning, debugging, and testing PART 4 BEYOND MAPREDUCE SQL on Hadoop Writing a YARN application

---

## APACHE SPARK IN 24 HOURS, SAMS TEACH YOURSELF

---

[Sams Publishing](#) Apache Spark is a fast, scalable, and flexible open source distributed processing engine for big data systems and is one of the most active open source big data projects to date. In just 24 lessons of one hour or less, Sams Teach Yourself Apache Spark in 24 Hours helps you build practical Big Data solutions that leverage Spark's amazing speed, scalability, simplicity, and versatility. This book's straightforward, step-by-step approach shows you how to deploy, program, optimize, manage, integrate, and extend Spark-now, and for years to come. You'll discover how to create powerful solutions encompassing cloud computing, real-time stream processing, machine learning, and more. Every lesson builds on what you've already learned, giving you a rock-solid foundation for real-world success. Whether you are a data analyst, data engineer, data scientist, or data steward, learning Spark will help you to advance your career or embark on a new career in the booming area of Big Data. Learn how to • Discover what Apache Spark does and how it fits into the Big Data landscape • Deploy and run Spark locally or in the cloud • Interact with Spark from the shell • Make the most of the Spark Cluster Architecture • Develop Spark applications with Scala and functional Python • Program with the Spark API, including transformations and actions • Apply practical data engineering/analysis approaches designed for Spark • Use Resilient Distributed Datasets (RDDs) for caching, persistence, and output • Optimize Spark solution performance • Use Spark with SQL (via Spark SQL) and with NoSQL (via Cassandra) • Leverage cutting-edge functional programming techniques • Extend Spark with streaming, R, and Sparkling Water • Start building Spark-based machine learning and graph-processing applications • Explore advanced messaging technologies, including Kafka • Preview and prepare for Spark's next generation of innovations Instructions walk you through common questions, issues, and tasks; Q-and-As, Quizzes, and Exercises build and test your knowledge; "Did You Know?" tips offer insider advice and shortcuts; and "Watch Out!" alerts help you avoid pitfalls. By the time you're finished, you'll be comfortable using Apache Spark to solve a wide spectrum of Big Data problems.

---

## MAKING APPS WITH MOQUI

---



---

### HOLISTIC ENTERPRISE APPLICATIONS MADE EASY

---

Making Apps with Moqui is the official documentation for Moqui Framework and includes a comprehensive summary of Mantle Business Artifacts. Starting with basic concepts and a tutorial to try things right away, it builds to complete examples of end-to-end business processes including procure to pay, order to cash, and work plan to cash. The framework topics cover data and service tier tools, user and system interfaces, security, and performance. With dozens of diagrams and screen shots, and thousands of lines of code and configuration examples, this book gives you ideas of what you can do with Moqui Framework and shows you how too. This includes things as simple as defining your data model with entities to more advanced things like building hierarchical data documents based on entity data and feeding them to other systems or indexing and searching the documents through simple configuration. Learn how to easily build remote and local services that handle validation, security, transaction management, and much more. Build screens quickly with a wide variety of dynamic widgets and forms styled any way you wish, or even define your own widgets to use consistently across your applications. Handle large scale and multi-tenant systems. Track your application use and performance. Implicitly handle multiple languages, currencies and other localization details. Control access to resources across all tiers through flexible authc and authz configuration. Written by the founder of Moqui and Mantle, and an enterprise application architect with 15 years of open source and commercial experience, this book provides the most accurate and useful information available for building modern enterprise applications with some of the best open source tools and technologies.

---

## HEAD FIRST SERVLETS AND JSP

---



---

### PASSING THE SUN CERTIFIED WEB COMPONENT DEVELOPER EXAM

---

["O'Reilly Media, Inc."](#) Looking to study up for the new J2EE 1.5 Sun Certified Web Component Developer (SCWCD) exam? This book will get you way up to speed on the technology you'll know it so well, in fact, that you can pass the brand new J2EE 1.5 exam. If that's what you want to do, that is. Maybe you don't care about the exam, but need to use servlets and JSPs in your next project. You're working on a deadline. You're over the legal limit for caffeine. You can't waste your time with a book that makes sense only AFTER you're an expert (or worse, one that puts you to sleep). Learn how to write servlets and JSPs, what makes a web container tick (and what ticks it off), how to use JSP's Expression Language (EL for short), and how to write deployment descriptors for your web applications. Master the c:out tag, and get a handle on exactly what's changed since the older J2EE 1.4 exam. You don't just pass the new J2EE 1.5 SCWCD exam, you'll understand this stuff and put it to work immediately. Head First Servlets and JSP doesn't just give you a bunch of facts to memorize; it drives knowledge straight into your brain. You'll interact with servlets and JSPs in ways that help you learn quickly and deeply. And when you're through with the book, you can take a brand-new mock exam, created specifically to simulate the real test-taking experience.

---

## DATABASE ADMINISTRATION

---



---

### THE COMPLETE GUIDE TO PRACTICES AND PROCEDURES

---

[Addison-Wesley Professional](#) A thorough reference on database administration outlines a variety of DBA roles and responsibilities and discusses such topics as data modeling and normalization, database/application design, change

management, database security and data integrity, performance issues, disaster planning, and other essentials. Original. (Advanced)

---

## PRO HADOOP DATA ANALYTICS

---

---

### DESIGNING AND BUILDING BIG DATA SYSTEMS USING THE HADOOP ECOSYSTEM

---

**Apress** Learn advanced analytical techniques and leverage existing tool kits to make your analytic applications more powerful, precise, and efficient. This book provides the right combination of architecture, design, and implementation information to create analytical systems that go beyond the basics of classification, clustering, and recommendation. Pro Hadoop Data Analytics emphasizes best practices to ensure coherent, efficient development. A complete example system will be developed using standard third-party components that consist of the tool kits, libraries, visualization and reporting code, as well as support glue to provide a working and extensible end-to-end system. The book also highlights the importance of end-to-end, flexible, configurable, high-performance data pipeline systems with analytical components as well as appropriate visualization results. You'll discover the importance of mix-and-match or hybrid systems, using different analytical components in one application. This hybrid approach will be prominent in the examples. **What You'll Learn** Build big data analytic systems with the Hadoop ecosystem Use libraries, tool kits, and algorithms to make development easier and more effective Apply metrics to measure performance and efficiency of components and systems Connect to standard relational databases, noSQL data sources, and more Follow case studies with example components to create your own systems **Who This Book Is For** Software engineers, architects, and data scientists with an interest in the design and implementation of big data analytical systems using Hadoop, the Hadoop ecosystem, and other associated technologies.

---

## IBM PLATFORM COMPUTING SOLUTIONS

---

**IBM Redbooks** This IBM® Platform Computing Solutions Redbooks® publication is the first book to describe each of the available offerings that are part of the IBM portfolio of Cloud, analytics, and High Performance Computing (HPC) solutions for our clients. This IBM Redbooks publication delivers descriptions of the available offerings from IBM Platform Computing that address challenges for our clients in each industry. We include a few implementation and testing scenarios with selected solutions. This publication helps strengthen the position of IBM Platform Computing solutions with a well-defined and documented deployment model within an IBM System x® environment. This deployment model offers clients a planned foundation for dynamic cloud infrastructure, provisioning, large-scale parallel HPC application development, cluster management, and grid applications. This IBM publication is targeted to IT specialists, IT architects, support personnel, and clients. This book is intended for anyone who wants information about how IBM Platform Computing solutions use IBM to provide a wide array of client solutions.

---

## SERVLET & JSP: A BEGINNER'S TUTORIAL

---

**Brainy Software Inc** Servlet and JavaServer Pages (JSP) are the underlying technologies for developing web applications in Java. They are essential for any programmer to master in order to effectively use frameworks such as JavaServer Faces, Struts 2, or Spring MVC. Covering Servlet 3.1 and JSP 2.3, this book explains the important programming concepts and design models in Java web development as well as related technologies and new features in the latest versions of Servlet and JSP. With comprehensive coverage and a lot of examples, this book is a guide to building real-world applications.

---

## ARCHITECTING MODERN DATA PLATFORMS

---

---

### A GUIDE TO ENTERPRISE HADOOP AT SCALE

---

**"O'Reilly Media, Inc."** There's a lot of information about big data technologies, but splicing these technologies into an end-to-end enterprise data platform is a daunting task not widely covered. With this practical book, you'll learn how to build big data infrastructure both on-premises and in the cloud and successfully architect a modern data platform. Ideal for enterprise architects, IT managers, application architects, and data engineers, this book shows you how to overcome the many challenges that emerge during Hadoop projects. You'll explore the vast landscape of tools available in the Hadoop and big data realm in a thorough technical primer before diving into: **Infrastructure:** Look at all component layers in a modern data platform, from the server to the data center, to establish a solid foundation for data in your enterprise **Platform:** Understand aspects of deployment, operation, security, high availability, and disaster recovery, along with everything you need to know to integrate your platform with the rest of your enterprise IT **Taking Hadoop to the cloud:** Learn the important architectural aspects of running a big data platform in the cloud while maintaining enterprise security and high availability

---

## IBM SOFTWARE FOR SAP SOLUTIONS

---

**IBM Redbooks** SAP is a market leader in enterprise business application software. SAP solutions provide a rich set of composable application modules, and configurable functional capabilities that are expected from a comprehensive enterprise business application software suite. In most cases, companies that adopt SAP software remain heterogeneous enterprises running both SAP and non-SAP systems to support their business processes. Regardless of the specific scenario, in heterogeneous enterprises most SAP implementations must be integrated with a variety of non-SAP enterprise systems: **Portals** **Messaging infrastructure** **Business process management (BPM) tools** **Enterprise**

Content Management (ECM) methods and tools Business analytics (BA) and business intelligence (BI) technologies Security Systems of record Systems of engagement The tooling included with SAP software addresses many needs for creating SAP-centric environments. However, the classic approach to implementing SAP functionality generally leaves the business with a rigid solution that is difficult and expensive to change and enhance. When SAP software is used in a large, heterogeneous enterprise environment, SAP clients face the dilemma of selecting the correct set of tools and platforms to implement SAP functionality, and to integrate the SAP solutions with non-SAP systems. This IBM® Redbooks® publication explains the value of integrating IBM software with SAP solutions. It describes how to enhance and extend pre-built capabilities in SAP software with best-in-class IBM enterprise software, enabling clients to maximize return on investment (ROI) in their SAP investment and achieve a balanced enterprise architecture approach. This book describes IBM Reference Architecture for SAP, a prescriptive blueprint for using IBM software in SAP solutions. The reference architecture is focused on defining the use of IBM software with SAP, and is not intended to address the internal aspects of SAP components. The chapters of this book provide a specific reference architecture for many of the architectural domains that are each important for a large enterprise to establish common strategy, efficiency, and balance. The majority of the most important architectural domain topics, such as integration, process optimization, master data management, mobile access, Enterprise Content Management, business intelligence, DevOps, security, systems monitoring, and so on, are covered in the book. However, there are several other architectural domains which are not included in the book. This is not to imply that these other architectural domains are not important or are less important, or that IBM does not offer a solution to address them. It is only reflective of time constraints, available resources, and the complexity of assembling a book on an extremely broad topic. Although more content could have been added, the authors feel confident that the scope of architectural material that has been included should provide organizations with a fantastic head start in defining their own enterprise reference architecture for many of the important architectural domains, and it is hoped that this book provides great value to those reading it. This IBM Redbooks publication is targeted to the following audiences: Client decision makers and solution architects leading enterprise transformation projects and wanting to gain further insight so that they can benefit from the integration of IBM software in large-scale SAP projects. IT architects and consultants integrating IBM technology with SAP solutions.

---

## **DATABASE DESIGN AND IMPLEMENTATION**

---

### **SECOND EDITION**

---

[Springer Nature](#) This textbook examines database systems from the viewpoint of a software developer. This perspective makes it possible to investigate why database systems are the way they are. It is of course important to be able to write queries, but it is equally important to know how they are processed. We e.g. don't want to just use JDBC; we also want to know why the API contains the classes and methods that it does. We need a sense of how hard is it to write a disk cache or logging facility. And what exactly is a database driver, anyway? The first two chapters provide a brief overview of database systems and their use. Chapter 1 discusses the purpose and features of a database system and introduces the Derby and SimpleDB systems. Chapter 2 explains how to write a database application using Java. It presents the basics of JDBC, which is the fundamental API for Java programs that interact with a database. In turn, Chapters 3-11 examine the internals of a typical database engine. Each chapter covers a different database component, starting with the lowest level of abstraction (the disk and file manager) and ending with the highest (the JDBC client interface); further, the respective chapter explains the main issues concerning the component, and considers possible design decisions. As a result, the reader can see exactly what services each component provides and how it interacts with the other components in the system. By the end of this part, s/he will have witnessed the gradual development of a simple but completely functional system. The remaining four chapters then focus on efficient query processing, and focus on the sophisticated techniques and algorithms that can replace the simple design choices described earlier. Topics include indexing, sorting, intelligent buffer usage, and query optimization. This text is intended for upper-level undergraduate or beginning graduate courses in Computer Science. It assumes that the reader is comfortable with basic Java programming; advanced Java concepts (such as RMI and JDBC) are fully explained in the text. The respective chapters are complemented by "end-of-chapter readings" that discuss interesting ideas and research directions that went unmentioned in the text, and provide references to relevant web pages, research articles, reference manuals, and books. Conceptual and programming exercises are also included at the end of each chapter. Students can apply their conceptual knowledge by examining the SimpleDB (a simple but fully functional database system created by the author and provided online) code and modifying it.

---

## **DESIGN PATTERNS**

---

### **ELEMENTS OF REUSABLE OBJECT-ORIENTED SOFTWARE**

---

[Pearson Deutschland GmbH](#) Software -- Software Engineering.

---

### **MOBILE CLOUD COMPUTING**

---

### **MODELS, IMPLEMENTATION, AND SECURITY**

---

[CRC Press](#) **Mobile Cloud Computing: Models, Implementation, and Security** provides a comprehensive introduction to mobile cloud computing, including key concepts, models, and relevant applications. The book focuses on novel and advanced algorithms, as well as mobile app development. The book begins with an overview of mobile cloud computing

concepts, models, and service deployments, as well as specific cloud service models. It continues with the basic mechanisms and principles of mobile computing, as well as virtualization techniques. The book also introduces mobile cloud computing architecture, design, key techniques, and challenges. The second part of the book covers optimizations of data processing and storage in mobile clouds, including performance and green clouds. The crucial optimization algorithm in mobile cloud computing is also explored, along with big data and service computing. Security issues in mobile cloud computing are covered in-depth, including a brief introduction to security and privacy issues and threats, as well as privacy protection techniques in mobile systems. The last part of the book features the integration of service-oriented architecture with mobile cloud computing. It discusses web service specifications related to implementations of mobile cloud computing. The book not only presents critical concepts in mobile cloud systems, but also drives readers to deeper research, through open discussion questions. Practical case studies are also included. Suitable for graduate students and professionals, this book provides a detailed and timely overview of mobile cloud computing for a broad range of readers.

---

## A PRACTICAL GUIDE TO CONTINUOUS DELIVERY

---

[Addison-Wesley Professional](#) **Using Continuous Delivery**, you can bring software into production more rapidly, with greater reliability. **A Practical Guide to Continuous Delivery** is a 100% practical guide to building Continuous Delivery pipelines that automate rollouts, improve reproducibility, and dramatically reduce risk. Eberhard Wolff introduces a proven Continuous Delivery technology stack, including Docker, Chef, Vagrant, Jenkins, Graphite, the ELK stack, JBehave, and Gatling. He guides you through applying these technologies throughout build, continuous integration, load testing, acceptance testing, and monitoring. Wolff's start-to-finish example projects offer the basis for your own experimentation, pilot programs, and full-fledged deployments. **A Practical Guide to Continuous Delivery** is for everyone who wants to introduce Continuous Delivery, with or without DevOps. For managers, it introduces core processes, requirements, benefits, and technical consequences. Developers, administrators, and architects will gain essential skills for implementing and managing pipelines, and for integrating Continuous Delivery smoothly into software architectures and IT organizations. Understand the problems that Continuous Delivery solves, and how it solves them Establish an infrastructure for maximum software automation Leverage virtualization and Platform as a Service (PAAS) cloud solutions Implement build automation and continuous integration with Gradle, Maven, and Jenkins Perform static code reviews with SonarQube and repositories to store build artifacts Establish automated GUI and textual acceptance testing with behavior-driven design Ensure appropriate performance via capacity testing Check new features and problems with exploratory testing Minimize risk throughout automated production software rollouts Gather and analyze metrics and logs with Elasticsearch, Logstash, Kibana (ELK), and Graphite Manage the introduction of Continuous Delivery into your enterprise Architect software to facilitate Continuous Delivery of new capabilities

---

## DROOLS DEVELOPER'S COOKBOOK

---

[Packt Publishing Ltd](#) Part of Packt's cookbook series, this book is packed with easy to follow recipes containing step-by-step instructions. The book is designed in such a way that you can read it chapter by chapter, or refer to the tasks in no particular order. This book is for Drools developers who want to improve their current working methods and discover new features to apply to their projects. Readers are expected to be familiar with the basics of the Drools platform as well as Java.

---

## CLIENT-SERVER WEB APPS WITH JAVASCRIPT AND JAVA

---

["O'Reilly Media, Inc."](#) As a Java programmer, how can you tackle the disruptive client-server approach to web development? With this comprehensive guide, you'll learn how today's client-side technologies and web APIs work with various Java tools. Author Casimir Saternos provides the big picture of client-server development, and then takes you through many practical client-server architectures. You'll work with hands-on projects in several chapters to get a feel for the topics discussed. User habits, technologies, and development methods have drastically altered web app design in recent years. But the Web itself hasn't changed. This book shows you how to build apps that conform to the web's underlying architecture. Learn the advantages of using separate client and server tiers, including code organization and speedy prototyping Explore the major tools, frameworks, and starter projects used in JavaScript development Dive into web API design and REST style of software architecture Understand Java's alternatives to traditional packaging methods and application server deployment Build projects with lightweight servers, using jQuery with Jython, and Sinatra with Angular Create client-server web apps with traditional Java web application servers and libraries

---

## DEPENDENCY INJECTION

---



---

## DESIGN PATTERNS USING SPRING AND GUICE

---

[Simon and Schuster](#) **Dependency Injection** is an in-depth guide to the current best practices focusing the **Dependency Injection** pattern-the key concept in Spring and the rapidly-growing Google Guice. It explores **Dependency Injection**, sometimes called **Inversion of Control**, in fine detail with numerous practical examples. Developers will learn to apply important techniques, focusing on their strengths and limitations, with a particular emphasis on pitfalls, corner-cases, and best practices. This book is written for developers and architects who want to understand **Dependency Injection** and successfully leverage popular DI technologies such as Spring, Google Guice, PicoContainer, and many others. The book explores many small examples of anchor concepts and unfolds a larger example to show the big picture. Written primarily from a Java point-of-view, this book is appropriate for any developer with a working knowledge of object-

oriented programming in Java, Ruby, or C#. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book.

---

## SPRING 5 RECIPES

---

### A PROBLEM-SOLUTION APPROACH

---

**Apress** Solve all your Spring 5 problems using complete and real-world code examples. When you start a new project, you'll be able to copy the code and configuration files from this book, and then modify them for your needs. This can save you a great deal of work over creating a project from scratch. The recipes in Spring 5 Recipes cover Spring fundamentals such as Spring IoC container, Spring AOP/ AspectJ, and more. Other recipes include Spring enterprise solutions for topics such as Spring Java EE integration, Spring Integration, Spring Batch, Spring Remoting, messaging, transactions, and working with big data and the cloud using Hadoop and MongoDB. Finally, Spring web recipes cover Spring MVC, other dynamic scripting, integration with the popular Grails Framework (and Groovy), REST/web services, and more. You'll also see recipes on new topics such as Spring Framework 5, reactive Spring, Spring 5 microservices, the functional web framework and much more. This book builds upon the best-selling success of the previous editions and focuses on the latest Spring Framework features for building enterprise Java applications. What You'll Learn Get re-usable code recipes and snippets for core Spring, annotations and other development tools Access Spring MVC for web development Work with Spring REST and microservices for web services development and integration into your enterprise Java applications Use Spring Batch, NoSQL and big data for building and integrating various cloud computing services and resources Integrate Java Enterprise Edition and other Java APIs for use in Spring Use Grails code and much more Who This Book Is For Experienced Java and Spring programmers.

---

### IMPLEMENTING AN IBM HIGH-PERFORMANCE COMPUTING SOLUTION ON IBM POWER SYSTEM S822LC

---

**IBM Redbooks** This IBM® Redbooks® publication demonstrates and documents that IBM Power Systems™ high-performance computing and technical computing solutions deliver faster time to value with powerful solutions. Configurable into highly scalable Linux clusters, Power Systems offer extreme performance for demanding workloads such as genomics, finance, computational chemistry, oil and gas exploration, and high-performance data analytics. This book delivers a high-performance computing solution implemented on the IBM Power System S822LC. The solution delivers high application performance and throughput based on its built-for-big-data architecture that incorporates IBM POWER8® processors, tightly coupled Field Programmable Gate Arrays (FPGAs) and accelerators, and faster I/O by using Coherent Accelerator Processor Interface (CAPI). This solution is ideal for clients that need more processing power while simultaneously increasing workload density and reducing datacenter floor space requirements. The Power S822LC offers a modular design to scale from a single rack to hundreds, simplicity of ordering, and a strong innovation roadmap for graphics processing units (GPUs). This publication is targeted toward technical professionals (consultants, technical support staff, IT Architects, and IT Specialists) responsible for delivering cost effective high-performance computing (HPC) solutions that help uncover insights from their data so they can optimize business results, product development, and scientific discoveries

---

## HADOOP OPERATIONS

---

### A GUIDE FOR DEVELOPERS AND ADMINISTRATORS

---

**"O'Reilly Media, Inc."** If you've been asked to maintain large and complex Hadoop clusters, this book is a must. Demand for operations-specific material has skyrocketed now that Hadoop is becoming the de facto standard for truly large-scale data processing in the data center. Eric Sammer, Principal Solution Architect at Cloudera, shows you the particulars of running Hadoop in production, from planning, installing, and configuring the system to providing ongoing maintenance. Rather than run through all possible scenarios, this pragmatic operations guide calls out what works, as demonstrated in critical deployments. Get a high-level overview of HDFS and MapReduce: why they exist and how they work Plan a Hadoop deployment, from hardware and OS selection to network requirements Learn setup and configuration details with a list of critical properties Manage resources by sharing a cluster across multiple groups Get a runbook of the most common cluster maintenance tasks Monitor Hadoop clusters—and learn troubleshooting with the help of real-world war stories Use basic tools and techniques to handle backup and catastrophic failure

---

## LUCENE IN ACTION

---

**Simon and Schuster** When Lucene first hit the scene five years ago, it was nothing short of amazing. By using this open-source, highly scalable, super-fast search engine, developers could integrate search into applications quickly and efficiently. A lot has changed since then—search has grown from a "nice-to-have" feature into an indispensable part of most enterprise applications. Lucene now powers search in diverse companies including Akamai, Netflix, LinkedIn, Technorati, HotJobs, Epiphany, FedEx, Mayo Clinic, MIT, New Scientist Magazine, and many others. Some things remain the same, though. Lucene still delivers high-performance search features in a disarmingly easy-to-use API. Due to its vibrant and diverse open-source community of developers and users, Lucene is relentlessly improving, with evolutions to APIs, significant new features such as payloads, and a huge increase (as much as 8x) in indexing speed with Lucene 2.3. And with clear writing, reusable examples, and unmatched advice on best practices, Lucene in Action, Second Edition is still the definitive guide to developing with Lucene. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book.

---

## PROGRAMMING HIVE

---

"O'Reilly Media, Inc." Describes the features and functions of Apache Hive, the data infrastructure for Hadoop.

---

## IBM TECHNICAL COMPUTING CLOUDS

---

[IBM Redbooks](#) This IBM® Redbooks® publication highlights IBM Technical Computing as a flexible infrastructure for clients looking to reduce capital and operational expenditures, optimize energy usage, or re-use the infrastructure. This book strengthens IBM SmartCloud® solutions, in particular IBM Technical Computing clouds, with a well-defined and documented deployment model within an IBM System x® or an IBM Flex System™. This provides clients with a cost-effective, highly scalable, robust solution with a planned foundation for scaling, capacity, resilience, optimization, automation, and monitoring. This book is targeted toward technical professionals (consultants, technical support staff, IT Architects, and IT Specialists) responsible for providing cloud-computing solutions and support.

---

## PROGRAMMING PIG

---

"O'Reilly Media, Inc." This guide is an ideal learning tool and reference for Apache Pig, the programming language that helps programmers describe and run large data projects on Hadoop. With Pig, they can analyze data without having to create a full-fledged application--making it easy for them to experiment with new data sets.

---

## APACHE KARAF COOKBOOK

---

[Packt Publishing Ltd](#) This book is intended for developers who have some familiarity with Apache Karaf and who want a quick reference for practical, proven tips on how to perform common tasks such as configuring Pax modules deployed in Apache Karaf, Extending HttpService with Apache Karaf. You should have working knowledge of Apache karaf, as the book provides a deeper understanding of the capabilities of Apache Karaf.

---

## CLOUD COMPUTING AND SOFTWARE SERVICES

---

### THEORY AND TECHNIQUES

---

[CRC Press](#) Whether you're already in the cloud, or determining whether or not it makes sense for your organization, *Cloud Computing and Software Services: Theory and Techniques* provides the technical understanding needed to develop and maintain state-of-the-art cloud computing and software services. From basic concepts and recent research findings to fut

---

## MASTERING CLOUD COMPUTING

---

### FOUNDATIONS AND APPLICATIONS PROGRAMMING

---

[Newnes](#) *Mastering Cloud Computing* is designed for undergraduate students learning to develop cloud computing applications. Tomorrow's applications won't live on a single computer but will be deployed from and reside on a virtual server, accessible anywhere, any time. Tomorrow's application developers need to understand the requirements of building apps for these virtual systems, including concurrent programming, high-performance computing, and data-intensive systems. The book introduces the principles of distributed and parallel computing underlying cloud architectures and specifically focuses on virtualization, thread programming, task programming, and map-reduce programming. There are examples demonstrating all of these and more, with exercises and labs throughout. Explains how to make design choices and tradeoffs to consider when building applications to run in a virtual cloud environment Real-world case studies include scientific, business, and energy-efficiency considerations

---

## THINKING IN JAVA

---

[Prentice Hall Professional](#) An overview of the programming language's fundamentals covers syntax, initialization, implementation, classes, error handling, objects, applets, multiple threads, projects, and network programming.

---

## MAPREDUCE DESIGN PATTERNS

---

### BUILDING EFFECTIVE ALGORITHMS AND ANALYTICS FOR HADOOP AND OTHER SYSTEMS

---

"O'Reilly Media, Inc." Until now, design patterns for the MapReduce framework have been scattered among various research papers, blogs, and books. This handy guide brings together a unique collection of valuable MapReduce patterns that will save you time and effort regardless of the domain, language, or development framework you're using. Each pattern is explained in context, with pitfalls and caveats clearly identified to help you avoid common design mistakes when modeling your big data architecture. This book also provides a complete overview of MapReduce that explains its origins and implementations, and why design patterns are so important. All code examples are written for Hadoop. Summarization patterns: get a top-level view by summarizing and grouping data Filtering patterns: view data subsets such as records generated from one user Data organization patterns: reorganize data to work with other systems, or to make MapReduce analysis easier Join patterns: analyze different datasets together to discover interesting relationships Metapatterns: piece together several patterns to solve multi-stage problems, or to perform several analytics in the same job Input and output patterns: customize the way you use Hadoop to load or store data "A clear exposition of MapReduce programs for common data processing patterns—this book is indispensable for

---

anyone using Hadoop." --Tom White, author of Hadoop: The Definitive Guide

---

## INTERNET OF THINGS

---

### CONVERGING TECHNOLOGIES FOR SMART ENVIRONMENTS AND INTEGRATED ECOSYSTEMS

---

River Publishers The book aims to provide a broad overview of various topics of the Internet of Things (IoT) from the research and development priorities to enabling technologies, architecture, security, privacy, interoperability and industrial applications. It is intended to be a stand-alone book in a series that covers the Internet of Things activities of the IERC - Internet of Things European Research Cluster - from technology to international cooperation and the global "state of play." The book builds on the ideas put forward by the European Research Cluster on the Internet of Things Strategic Research and Innovation Agenda and presents views and state of the art results on the challenges facing the research, development and deployment of IoT at the global level. Today we see the integration of Industrial, Business and Consumer Internet which is bringing together the Internet of People, Internet of Things, Internet of Energy, Internet of Vehicles, Internet of Media, Services and Enterprises in forming the backbone of the digital economy, the digital society and the foundation for the future knowledge and innovation based economy. These developments are supporting solutions for the emerging challenges of public health, aging population, environmental protection and climate change, the conservation of energy and scarce materials, enhancements to safety and security and the continuation and growth of economic prosperity. Penetration of smartphones and advances in nanoelectronics, cyber-physical systems, wireless communication, software, and Cloud computing technology will be the main drivers for IoT development. The IoT contribution is seen in the increased value of information created by the number of interconnections among things and the transformation of the processed information into knowledge shared into the Internet of Everything. The connected devices are part of ecosystems connecting people, processes, data, and things which are communicating in the Cloud using the increased storage and computing power while attempting to standardize communication and metadata. In this context, the next generation of Cloud computing technologies will need to be flexible enough to scale autonomously, adaptive enough to handle constantly changing connections and resilient enough to stand up to the huge flows of data that will occur. In 2025, analysts forecast that there will be six devices per human on the planet, which means around 50 billion more connected devices over the next 12 years. The Internet of Things market is connected to this anticipated device growth from industrial Machine to Machine (M2M) systems, smart meters and wireless sensors. Internet of Things technology will generate new services and new interfaces by creating smart environments and smart spaces with applications ranging from Smart Cities, Smart Transport, Buildings, Energy, Grid, to Smart Health and Life.