Data Access For Highly Scalable Solutions Using Sql Nosql And Polyglot Persistence Microsoft Patterns Practices

Thank you categorically much for downloading **Data Access For Highly Scalable Solutions Using Sql Nosql And Polyglot Persistence Microsoft Patterns Practices**.Most likely you have knowledge that, people have look numerous time for their favorite books next this Data Access For Highly Scalable Solutions Using Sql Nosql And Polyglot Persistence Microsoft Patterns Practices, but end happening in harmful downloads.

Rather than enjoying a good book past a cup of coffee in the afternoon, instead they juggled behind some harmful virus inside their computer. **Data Access For Highly Scalable Solutions Using Sql Nosql And Polyglot Persistence Microsoft Patterns Practices** is reachable in our digital library an online permission to it is set as public thus you can download it instantly. Our digital library saves in multipart countries, allowing you to get the most less latency period to download any of our books once this one. Merely said, the Data Access For Highly Scalable Solutions Using Sql Nosql And Polyglot Persistence Microsoft Patterns Practices is universally compatible in imitation of any devices to read.

MongoDB Applied Design Patterns - Rick Copeland 2013-03-04 Whether you're building a social media site or an internal-use enterprise application, this hands-on guide shows you the connection between MongoDB and the business problems it's designed to solve. You'll learn how to apply MongoDB design patterns to several challenging domains, such as ecommerce, content management, and online gaming. Using Python and JavaScript code examples, you'll discover how MongoDB lets you scale your data model while simplifying the development process. Many businesses launch NoSQL databases without understanding the techniques for using their features most effectively. This book demonstrates the benefits of document embedding, polymorphic schemas, and other MongoDB patterns for tackling specific big data use cases, including: Operational intelligence: Perform real-time analytics of business data Ecommerce: Use MongoDB as a product catalog master or inventory management system Content management: Learn methods for storing content nodes, binary assets, and discussions Online advertising networks: Apply techniques for frequency capping ad impressions, and keyword targeting and bidding Social networking: Learn how to store a complex social graph, modeled after Google+ Online gaming: Provide concurrent access to character and world data for a multiplayer role-playing game

.NET Design Patterns - Praseed Pai 2017-01-31 Explore the world of .NET design patterns and bring the benefits that the right patterns can offer to your toolkit today About This Book Dive into the powerful fundamentals of .NET

framework for software development The code is explained piece by piece and the application of the pattern is also showcased. This fast-paced guide shows you how to implement the patterns into your existing applications Who This Book Is For This book is for those with familiarity with .NET development who would like to take their skills to the next level and be in the driver's seat when it comes to modern development techniques. Basic object-oriented C# programming experience and an elementary familiarity with the .NET framework library is required. What You Will Learn Put patterns and pattern catalogs into the right perspective Apply patterns for software development under C#/.NET Use GoF and other patterns in real-life development scenarios Be able to enrich your design vocabulary and well articulate your design thoughts Leverage object/functional programming by mixing OOP and FP Understand the reactive programming model using Rx and RxJs Writing compositional code using C# LINQ constructs Be able to implement concurrent/parallel programming techniques using idioms under .NET Avoiding pitfalls when creating compositional, readable, and maintainable code using imperative, functional, and reactive code. In Detail Knowing about design patterns enables developers to improve their code base, promoting code reuse and making their design more robust. This book focuses on the practical aspects of programming in .NET. You will learn about some of the relevant design patterns (and their application) that are most widely used. We start with classic object-oriented programming (OOP) techniques, evaluate parallel programming and concurrency models, enhance implementations

by mixing OOP and functional programming, and finally to the reactive programming model where functional programming and OOP are used in synergy to write better code. Throughout this book, we'll show you how to deal with architecture/design techniques, GoF patterns, relevant patterns from other catalogs, functional programming, and reactive programming techniques. After reading this book, you will be able to convincingly leverage these design patterns (factory pattern, builder pattern, prototype pattern, adapter pattern, facade pattern, decorator pattern, observer pattern and so on) for your programs. You will also be able to write fluid functional code in .NET that would leverage concurrency and parallelism! Style and approach This tutorial-based book takes a step-by-step approach. It covers the major patterns and explains them in a detailed manned along with code examples. *Docker and Kubernetes for Java Developers* - Jaroslaw Krochmalski 2017-08-30 Leverage the lethal combination of Docker and Kubernetes to automate deployment and management of Java applications About This Book Master using Docker and Kubernetes to build, deploy and manage Java applications in a jiff Learn how to create your own Docker image and customize your own cluster using Kubernetes Empower the journey from development to production using this practical guide. Who This Book Is For The book is aimed at Java developers who are eager to build, deploy, and manage applications very quickly using container technology. They need have no knowledge of Docker and Kubernetes. What You Will Learn Package Java applications into Docker images Understand the running of containers locally

Explore development and deployment options with Docker Integrate Docker into Maven builds Manage and monitor Java applications running on Kubernetes clusters Create Continuous Delivery pipelines for Java applications deployed to Kubernetes In Detail Imagine creating and testing Java EE applications on Apache Tomcat Server or Wildfly Application server in minutes along with deploying and managing Java applications swiftly. Sounds too good to be true? But you have a reason to cheer as such scenarios are only possible by leveraging Docker and Kubernetes. This book will start by introducing Docker and delve deep into its networking and persistent storage concepts. You will then proceed to learn how to refactor monolith application into separate services by building an application and then packaging it into Docker containers. Next, you will create an image containing Java Enterprise Application and later run it using Docker. Moving on, the book will focus on Kubernetes and its features and you will learn to deploy a Java application to Kubernetes using Maven and monitor a Java application in production. By the end of the book, you will get hands-on with some more advanced topics to further extend your knowledge about Docker and Kubernetes. Style and approach An easy-to-follow, practical guide that will help Java developers develop, deploy, and manage Java applications efficiently.
Heterogeneous Data Management, Polystores, and Analytics for Healthcare - Vijay Gadepally 2019-03-21 This book constitutes the refereed post-conference proceedings of International Workshops: Polystores and other Systems for Heterogeneous Data, Poly 2018, and Data

Management and Analytics for Medicine and Healthcare, DMAH 2018, in Rio de Janeiro, Brazil, in August 2018, held in conjunction with the 44th International Conference on Very Large Data Bases, VLDB 2018. The 11 regular papers presented together with 2 invited papers and 1 abstract of a keynote talk were carefully reviewed and selected from 16 initial submissions. The Poly 2018 Workshop focus on growing a larger and more diverse research agenda around data system solutions for heterogeneous data. The DMAH 2018 Workshop aims to foster exchange of information and discussions on innovative data management and analytics technologies.

Designing Platform Independent Mobile Apps and Services - Rocky Heckman 2016-09-20 Presents strategies to designing platform agnostic mobile apps connected to cloud based services that can handle heavy loads of modern computing Provides development patterns for platform agnostic app development and technologies Includes recommended standards and structures for easy adoption Covers portable and modular back-end architectures to support service agility and rapid development

Embracing Microservices Design - Ovais Mehboob Ahmed Khan 2021-10-29 Develop microservice-based enterprise applications with expert guidance to avoid failures and technological debt with the help of real-world examples Key FeaturesImplement the right microservices adoption strategy to transition from monoliths to microservicesExplore real-world use cases that explain anti-patterns and alternative practices in microservices developmentDiscover proven recommendations for avoiding architectural mistakes when designing microservicesBook

Description Microservices have been widely adopted for designing distributed enterprise apps that are flexible, robust, and fine-grained into services that are independent of each other. There has been a paradigm shift where organizations are now either building new apps on microservices or transforming existing monolithic apps into microservices-based architecture. This book explores the importance of anti-patterns and the need to address flaws in them with alternative practices and patterns. You'll identify common mistakes caused by a lack of understanding when implementing microservices and cover topics such as organizational readiness to adopt microservices, domain-driven design, and resiliency and scalability of microservices. The book further demonstrates the anti-patterns involved in re-platforming brownfield apps and designing distributed data architecture. You'll also focus on how to avoid communication and deployment pitfalls and understand cross-cutting concerns such as logging, monitoring, and security. Finally, you'll explore testing pitfalls and establish a framework to address isolation, autonomy, and standardization. By the end of this book, you'll have understood critical mistakes to avoid while building microservices and the right practices to adopt early in the product life cycle to ensure the success of a microservices initiative. What you will learnDiscover the responsibilities of different individuals involved in a microservices initiativeAvoid the common mistakes in architecting microservices for scalability and resiliencyUnderstand the importance of domain-driven design when developing microservicesIdentify the common pitfalls involved in

migrating monolithic applications to microservicesExplore communication strategies, along with their potential drawbacks and alternativesDiscover the importance of adopting governance, security, and monitoringUnderstand the role of CI/CD and testingWho this book is for This practical microservices book is for software architects, solution architects, and developers involved in designing microservices architecture and its development, who want to gain insights into avoiding pitfalls and drawbacks in distributed applications, and save time and money that might otherwise get wasted if microservices designs fail. Working knowledge of microservices is assumed to get the most out of this book.

Practical Microservices Architectural Patterns - Binildas Christudas 2019-06-25 Take your distributed applications to the next level and see what the reference architectures associated with microservices can do for you. This book begins by showing you the distributed computing architecture landscape and provides an in-depth view of microservices architecture. Following this, you will work with CQRS, an essential pattern for microservices, and get a view of how distributed messaging works. Moving on, you will take a deep dive into Spring Boot and Spring Cloud. Coming back to CQRS, you will learn how event-driven microservices work with this pattern, using the Axon 2 framework. This takes you on to how transactions work with microservices followed by advanced architectures to address non-functional aspects such as high availability and scalability. In the concluding part of the book you develop your own enterprise-grade microservices application

using the Axon framework and true BASE transactions, while making it as secure as possible. What You Will Learn Shift from monolith architecture to microservices Work with distributed and ACID transactionsBuild solid architectures without two-phase commit transactions Discover the high availability principles in microservices Who This Book Is For Java developers with basic knowledge of distributed and multi-threaded application architecture, and no knowledge of Spring Boot or Spring Cloud. Knowledge of CQRS and event-driven architecture is not mandatory as this book will cover these in depth.
Scalability Patterns - Chander Dhall 2018-07-20 In this book, the CEO of Cazton, Inc. and internationally-acclaimed speaker, Chander Dhall, demonstrates current website design scalability patterns and takes a pragmatic approach to explaining their pros and cons to show you how to select the appropriate pattern for your site. He then tests the patterns by deliberately forcing them to fail and exposing potential flaws before discussing how to design the optimal pattern to match your scale requirements. The author explains the use of polyglot programming and how to match the right patterns to your business needs. He also details several No-SQL patterns and explains the fundamentals of different paradigms of No-SQL by showing complementary strategies of using them along with relational databases to achieve the best results. He also teaches how to make the scalability pattern work with a real-world microservices pattern. With the proliferation of countless electronic devices and the ever growing number of Internet users, the scalability of websites has become an increasingly

important challenge. Scalability, even though highly coveted, may not be so easy to achieve. Think that you can't attain responsiveness along with scalability? Chander Dhall will demonstrate that, in fact, they go hand in hand. What You'll Learn Architect and develop applications so that they are easy to scale. Learn different scaling and partitioning options and the combinations. Learn techniques to speed up responsiveness. Deep dive into caching, column-family databases, document databases, search engines and RDBMS. Learn scalability and responsiveness concepts that are usually ignored. Effectively balance scalability, performance, responsiveness, and availability while minimizing downtime. Who This Book Is For Executives (CXOs), software architects , developers, and IT Pros Professional NoSQL - Shashank Tiwari 2011-08-31

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.
Enterprise Cloud Strategy - Barry Briggs 2016-01-07 How do you start? How should you build a plan for cloud migration for your entire portfolio? How will your organization be affected by these changes?

This book, based on real-world cloud experiences by enterprise IT teams, seeks to provide the answers to these questions. Here, you'll see what makes the cloud so compelling to enterprises; with which applications you should start your cloud journey; how your organization will change, and how skill sets will evolve; how to measure progress; how to think about security, compliance, and business buy-in; and how to exploit the ever-growing feature set that the cloud offers to gain strategic and competitive advantage. **Big Data For Dummies** - Judith S. Hurwitz 2013-04-02 Find the right big data solution for your business or organization Big data management is one of the major challenges facing business, industry, and not-for-profit organizations. Data sets such as customer transactions for a mega-retailer, weather patterns monitored by meteorologists, or social

network activity can quickly outpace the capacity of traditional data management tools. If you need to develop or manage big data solutions, you'll appreciate how these four experts define, explain, and guide you through this new and often confusing concept. You'll learn what it is, why it matters, and how to choose and implement solutions that work. Effectively managing big data is an issue of growing importance to businesses, not-for-profit organizations, government, and IT professionals Authors are experts in information management, big data, and a variety of solutions Explains big data in detail and discusses how to select and implement a solution, security concerns to consider, data storage and presentation issues, analytics, and much more Provides essential information in a no-nonsense, easy-to-understand style that is empowering Big Data For Dummies cuts through the confusion and helps you take charge of big data solutions for your organization.

SQL & NoSQL Databases - Andreas Meier 2019-07-05 This book offers a comprehensive introduction to relational (SQL) and non-relational (NoSQL) databases. The authors thoroughly review the current state of database tools and techniques, and examine coming innovations. The book opens with a broad look at data management, including an overview of information systems and databases, and an explanation of contemporary database types: SQL and NoSQL databases, and their respective management systems The nature and uses of Big Data A high-level view of the organization of data management Data Modeling and Consistency Chapter-length treatment is afforded Data Modeling in

both relational and graph databases, including enterprise-wide data architecture, and formulas for database design. Coverage of languages extends from an overview of operators, to SQL and and QBE (Query by Example), to integrity constraints and more. A full chapter probes the challenges of Ensuring Data Consistency, covering: Multi-User Operation Troubleshooting Consistency in Massive Distributed Data Comparison of the ACID and BASE consistency models, and more System Architecture also gets from its own chapter, which explores Processing of Homogeneous and Heterogeneous Data; Storage and Access Structures; Multi-dimensional Data Structures and Parallel Processing with MapReduce, among other topics. Post-Relational and NoSQL Databases The chapter on post-relational databases discusses the limits of SQL – and what lies beyond, including Multi-Dimensional Databases, Knowledge Bases and and Fuzzy Databases. A final chapter covers NoSQL Databases, along with Development of Non-Relational Technologies, Key-Value, Column-Family and Document Stores XML Databases and Graphic Databases, and more The book includes more than 100 tables, examples and illustrations, and each chapter offers a list of resources for further reading. SQL & NoSQL Databases conveys the strengths and weaknesses of relational and non-relational approaches, and shows how to undertake development for big data applications. The book benefits readers including students and practitioners working across the broad field of applied information technology. This textbook has been recommended and developed for university courses in Germany, Austria and Switzerland.

Seven Databases in Seven Weeks - Luc Perkins 2018-04-05 Data is getting bigger and more complex by the day, and so are your choices in handling it. Explore some of the most cutting-edge databases available - from a traditional relational database to newer NoSQL approaches - and make informed decisions about challenging data storage problems. This is the only comprehensive guide to the world of NoSQL databases, with in-depth practical and conceptual introductions to seven different technologies: Redis, Neo4J, CouchDB, MongoDB, HBase, Postgres, and DynamoDB. This second edition includes a new chapter on DynamoDB and updated content for each chapter. While relational databases such as MySQL remain as relevant as ever, the alternative, NoSQL paradigm has opened up new horizons in performance and scalability and changed the way we approach data-centric problems. This book presents the essential concepts behind each database alongside hands-on examples that make each technology come alive. With each database, tackle a real-world problem that highlights the concepts and features that make it shine. Along the way, explore five database models - relational, key/value, columnar, document, and graph - from the perspective of challenges faced by real applications. Learn how MongoDB and CouchDB are strikingly different, make your applications faster with Redis and more connected with Neo4J, build a cluster of HBase servers using cloud services such as Amazon's Elastic MapReduce, and more. This new edition brings a brand new chapter on DynamoDB, updated code samples and exercises, and a more up-to-date account of each database's feature set. Whether you're a programmer building the

next big thing, a data scientist seeking solutions to thorny problems, or a technology enthusiast venturing into new territory, you will find something to inspire you in this book. What You Need: You'll need a *nix shell (Mac OS or Linux preferred, Windows users will need Cygwin), Java 6 (or greater), and Ruby 1.8.7 (or greater). Each chapter will list the downloads required for that database.
Data Lake for Enterprises - Tomcy John 2017-05-31 A practical guide to implementing your enterprise data lake using Lambda Architecture as the base About This Book Build a full-fledged data lake for your organization with popular big data technologies using the Lambda architecture as the base Delve into the big data technologies required to meet modern day business strategies A highly practical guide to implementing enterprise data lakes with lots of examples and real-world use-cases Who This Book Is For Java developers and architects who would like to implement a data lake for their enterprise will find this book useful. If you want to get hands-on experience with the Lambda Architecture and big data technologies by implementing a practical solution using these technologies, this book will also help you. What You Will Learn Build an enterprise-level data lake using the relevant big data technologies Understand the core of the Lambda architecture and how to apply it in an enterprise Learn the technical details around Sqoop and its functionalities Integrate Kafka with Hadoop components to acquire enterprise data Use flume with streaming technologies for stream-based processing Understand stream- based processing with reference to Apache Spark Streaming Incorporate Hadoop components and know the

advantages they provide for enterprise data lakes Build fast, streaming, and high-performance applications using ElasticSearch Make your data ingestion process consistent across various data formats with configurability Process your data to derive intelligence using machine learning algorithms In Detail The term "Data Lake" has recently emerged as a prominent term in the big data industry. Data scientists can make use of it in deriving meaningful insights that can be used by businesses to redefine or transform the way they operate. Lambda architecture is also emerging as one of the very eminent patterns in the big data landscape, as it not only helps to derive useful information from historical data but also correlates real-time data to enable business to take critical decisions. This book tries to bring these two important aspects — data lake and lambda architecture—together. This book is divided into three main sections. The first introduces you to the concept of data lakes, the importance of data lakes in enterprises, and getting you up-to-speed with the Lambda architecture. The second section delves into the principal components of building a data lake using the Lambda architecture. It introduces you to popular big data technologies such as Apache Hadoop, Spark, Sqoop, Flume, and ElasticSearch. The third section is a highly practical demonstration of putting it all together, and shows you how an enterprise data lake can be implemented, along with several real-world use-cases. It also shows you how other peripheral components can be added to the lake to make it more efficient. By the end of this book, you will be able to choose the right big data technologies using the lambda architectural patterns to build your

enterprise data lake. Style and approach The book takes a pragmatic approach, showing ways to leverage big data technologies and lambda architecture to build an enterprise-level data lake.

Data Access for Highly-Scalable Solutions - Douglas McMurtry 2013-09-30 All applications use data, and most applications also need to store this data somewhere. In the world of business solutions, this often meant creating a relational database. However, relational technology is not always the best solution to meet the increasingly complex data-processing requirements of modern business systems, especially when this processing involves storing and retrieving massive amounts of data. The advent of NoSQL databases has changed the way in which organizations have started to think about the way in which they structure their data. There is no standard definition of what a NoSQL database is other than they are all non-relational. They are less generalized than relational databases, but the driving force behind most NoSQL databases is focused efficiency and high scalability. The downside of NoSQL is that no single database is likely to be able to support the complete range of business requirements mandated by your applications. How do you select the most appropriate database to use, or should you remain with the relational model? A modern business application is not restricted to using a single data store, and an increasing number of solutions are now based on a polyglot architecture. The key to designing a successful application is to understand which databases best meet the needs of the various parts of the system, and how to combine these databases into a single, seamless solution. This guide helps you understand

these challenges and enables you to apply the principles of NoSQL databases and polyglot solutions in your own environment. To help illustrate how to build a polyglot solution, this guide presents a case study of a fictitious company faced with building a highly scalable web application capable of supporting many thousands of concurrent users.

Microservices and Containers - Parminder Singh Kocher 2018-03-16 Transition to Microservices and DevOps to Transform Your Software Development Effectiveness Thanks to the tech sector's latest game-changing innovations—the Internet of Things (IoT), software-enabled networking, and software as a service (SaaS), to name a few—there is now a seemingly insatiable demand for platforms and architectures that can improve the process of application development and deployment. In Microservices and Containers, longtime systems architect and engineering team leader Parminder Kocher analyzes two of the hottest new technology trends: microservices and containers. Together, as Kocher demonstrates, microservices and Docker containers can bring unprecedented agility and scalability to application development and deployment, especially in large, complex projects where speed is crucial but small errors can be disastrous. Learn how to leverage microservices and Docker to drive modular architectural design, on-demand scalability, application performance and reliability, time-to-market, code reuse, and exponential improvements in DevOps effectiveness. Kocher offers detailed guidance and a complete roadmap for transitioning from monolithic architectures, as well as an

in-depth case study that walks the reader through the migration of an enterprise-class SOA system. Understand how microservices enable you to organize applications into standalone components that are easier to manage, update, and scale Decide whether microservices and containers are worth your investment, and manage the organizational learning curve associated with them Apply best practices for interprocess communication among microservices Migrate monolithic systems in an orderly fashion Understand Docker containers, installation, and interfaces Network, orchestrate, and manage Docker containers effectively Use Docker to maximize scalability in microservices-based applications Apply your learning with an in-depth, hands-on case study Whether you are a software architect/developer or systems professional looking to move on from older approaches or a manager trying to maximize the business value of these technologies, Microservices and Containers will be an invaluable addition to your library. Register your product at informit.com/register for convenient access to downloads, updates, and/or corrections as they become available.

Building Cloud Apps with Microsoft Azure - Scott Guthrie 2014-07-18 This ebook walks you through a patterns-based approach to building real-world cloud solutions. The patterns apply to the development process as well as to architecture and coding practices. The content is based on a presentation developed by Scott Guthrie and delivered by him at the Norwegian Developers Conference (NDC) in June of 2013 (part 1, part 2), and at Microsoft Tech Ed Australia in September 2013 (part 1,

part 2). Many others updated and augmented the content while transitioning it from video to written form. Who should read this book Developers who are curious about developing for the cloud, are considering a move to the cloud, or are new to cloud development will find here a concise overview of the most important concepts and practices they need to know. The concepts are illustrated with concrete examples, and each chapter includes links to other resources that provide more in-depth information. The examples and the links to additional resources are for Microsoft frameworks and services, but the principles illustrated apply to other web development frameworks and cloud environments as well. Developers who are already developing for the cloud may find ideas here that will help make them more successful. Each chapter in the series can be read independently, so you can pick and choose topics that you're interested in. Anyone who watched Scott Guthrie's "Building Real World Cloud Apps with Windows Azure" presentation and wants more details and updated information will find that here. Assumptions This ebook expects that you have experience developing web applications by using Visual Studio and ASP.NET. Familiarity with C# would be helpful in places.
ASP.NET Core Application Development - James Chambers 2016-11-29 This is the eBook of the printed book and may not include any media, website access codes, or print supplements that may come packaged with the bound book. Through four complete sprints, this book takes you through every step needed to build brand new cross-platform web apps with ASP.NET Core, and make them available on the Internet. You won't just master Microsoft's

revolutionary open source ASP.NET Core technology: you'll learn how to integrate the immense power of MVC, Docker, Azure Web Apps, Visual Studio and Visual Studio Code, C#, JavaScript, TypeScript, and Entity Framework. Working through the authors' carefully designed sprints, you'll start with a blank canvas, move through software architecture and design, adjusting to user feedback, recovering from mistakes, builds, testing, deployment, maintenance, refactoring, and more. Along the way, you'll learn techniques for delivering state-of-the-art software to users more rapidly and repeatably than ever before. *Design Patterns for Cloud Native Applications* - Kasun Indrasiri 2021-05-17 With the immense cost savings and scalability the cloud provides, the rationale for building cloud native applications is no longer in question. The real issue is how. With this practical guide, developers will learn about the most commonly used design patterns for building cloud native applications using APIs, data, events, and streams in both greenfield and brownfield development. You'll learn how to incrementally design, develop, and deploy large and effective cloud native applications that you can manage and maintain at scale with minimal cost, time, and effort. Authors Kasun Indrasiri and Sriskandarajah Suhothayan highlight use cases that effectively demonstrate the challenges you might encounter at each step. Learn the fundamentals of cloud native applications Explore key cloud native communication, connectivity, and composition patterns Learn decentralized data management techniques Use event-driven architecture to build distributed and scalable cloud native applications

Explore the most commonly used patterns for API management and consumption Examine some of the tools and technologies you'll need for building cloud native systems

Microsoft Visual C# 2013 Step by Step - John Sharp 2013-11-15
Your hands-on guide to Visual C# fundamentals Expand your expertise—and teach yourself the fundamentals of Microsoft Visual C# 2013. If you have previous programming experience but are new to Visual C# 2013, this tutorial delivers the step-by-step guidance and coding exercises you need to master core topics and techniques. Discover how to: Create and debug C# applications in Visual Studio 2013 Work with variables, identifiers, statements, operators, and methods Create interfaces and define abstract classes Write robust code that can catch and handle exceptions Display and edit data using data binding with Microsoft ADO.NET Entity Framework Respond to user input and touchscreen gestures Handle events arising from multiple sources Develop your first Windows 8.1 apps

Big Data Application Architecture Q&A - Nitin Sawant 2014-01-24
Big Data Application Architecture Pattern Recipes provides an insight into heterogeneous infrastructures, databases, and visualization and analytics tools used for realizing the architectures of big data solutions. Its problem-solution approach helps in selecting the right architecture to solve the problem at hand. In the process of reading through these problems, you will learn harness the power of new big data opportunities which various enterprises use to attain real-time profits. Big Data Application Architecture Pattern Recipes answers one of the most critical questions of this time 'how do you select the best

end-to-end architecture to solve your big data problem?'. The book deals with various mission critical problems encountered by solution architects, consultants, and software architects while dealing with the myriad options available for implementing a typical solution, trying to extract insight from huge volumes of data in real–time and across multiple relational and non-relational data types for clients from industries like retail, telecommunication, banking, and insurance. The patterns in this book provide the strong architectural foundation required to launch your next big data application. The architectures for realizing these opportunities are based on relatively less expensive and heterogeneous infrastructures compared to the traditional monolithic and hugely expensive options that exist currently. This book describes and evaluates the benefits of heterogeneity which brings with it multiple options of solving the same problem, evaluation of trade-offs and validation of 'fitness-for-purpose' of the solution. *The Art of Scalability* - Martin L. Abbott 2015-05-23 The Comprehensive, Proven Approach to IT Scalability–Updated with New Strategies, Technologies, and Case Studies In The Art of Scalability, Second Edition, leading scalability consultants Martin L. Abbott and Michael T. Fisher cover everything you need to know to smoothly scale products and services for any requirement. This extensively revised edition reflects new technologies, strategies, and lessons, as well as new case studies from the authors' pioneering consulting practice, AKF Partners. Writing for technical and nontechnical decision-makers, Abbott and Fisher cover everything that impacts scalability, including

architecture, process, people, organization, and technology. Their insights and recommendations reflect more than thirty years of experience at companies ranging from eBay to Visa, and Salesforce.com to Apple. You'll find updated strategies for structuring organizations to maximize agility and scalability, as well as new insights into the cloud (IaaS/PaaS) transition, NoSQL, DevOps, business metrics, and more. Using this guide's tools and advice, you can systematically clear away obstacles to scalability–and achieve unprecedented IT and business performance. Coverage includes • Why scalability problems start with organizations and people, not technology, and what to do about it • Actionable lessons from real successes and failures • Staffing, structuring, and leading the agile, scalable organization • Scaling processes for hyper-growth environments • Architecting scalability: proprietary models for clarifying needs and making choices–including 15 key success principles • Emerging technologies and challenges: data cost, datacenter planning, cloud evolution, and customer-aligned monitoring • Measuring availability, capacity, load, and performance

Data Mining and Data Warehousing - Parteek Bhatia 2019-06-27 Written in lucid language, this valuable textbook brings together fundamental concepts of data mining and data warehousing in a single volume. Important topics including information theory, decision tree, Naïve Bayes classifier, distance metrics, partitioning clustering, associate mining, data marts and operational data store are discussed comprehensively. The textbook is written to cater to the needs of undergraduate students of

computer science, engineering and information technology for a course on data mining and data warehousing. The text simplifies the understanding of the concepts through exercises and practical examples. Chapters such as classification, associate mining and cluster analysis are discussed in detail with their practical implementation using Weka and R language data mining tools. Advanced topics including big data analytics, relational data models and NoSQL are discussed in detail. Pedagogical features including unsolved problems and multiple-choice questions are interspersed throughout the book for better understanding. **Microservices from Theory to Practice: Creating Applications in IBM Bluemix Using the Microservices Approach** - Shahir Daya 2016-04-04 Microservices is an architectural style in which large, complex software applications are composed of one or more smaller services. Each of these microservices focuses on completing one task that represents a small business capability. These microservices can be developed in any programming language. They communicate with each other using language-neutral protocols, such as Representational State Transfer (REST), or messaging applications, such as IBM® MQ Light. This IBM Redbooks® publication gives a broad understanding of this increasingly popular architectural style, and provides some real-life examples of how you can develop applications using the microservices approach with IBM BluemixTM. The source code for all of these sample scenarios can be found on GitHub (https://github.com/). The book also presents some case studies from IBM products. We explain the architectural decisions

made, our experiences, and lessons learned when redesigning these products using the microservices approach. Information technology (IT) professionals interested in learning about microservices and how to develop or redesign an application in Bluemix using microservices can benefit from this book.

Accelerating Modernization with Agile Integration - Adeline SE Chun 2020-07-01
The organization pursuing digital transformation must embrace new ways to use and deploy integration technologies, so they can move quickly in a manner appropriate to the goals of multicloud, decentralization, and microservices. The integration layer must transform to allow organizations to move boldly in building new customer experiences, rather than forcing models for architecture and development that pull away from maximizing the organization's productivity. Many organizations have started embracing agile application techniques, such as microservice architecture, and are now seeing the benefits of that shift. This approach complements and accelerates an enterprise's API strategy. Businesses should also seek to use this approach to modernize their existing integration and messaging infrastructure to achieve more effective ways to manage and operate their integration services in their private or public cloud. This IBM® Redbooks® publication explores the merits of what we refer to as agile integration; a container-based, decentralized, and microservice-aligned approach for integration solutions that meets the demands of agility, scalability, and resilience required by digital transformation. It also discusses how the IBM Cloud Pak for Integration marks a

significant leap forward in integration technology by embracing both a cloud-native approach and container technology to achieve the goals of agile integration. The target audiences for this book are cloud integration architects, IT specialists, and application developers.
Real-Time Big Data Analytics - Sumit Gupta 2016-02-26 Design, process, and analyze large sets of complex data in real time About This Book Get acquainted with transformations and database-level interactions, and ensure the reliability of messages processed using Storm Implement strategies to solve the challenges of real-time data processing Load datasets, build queries, and make recommendations using Spark SQL Who This Book Is For If you are a Big Data architect, developer, or a programmer who wants to develop applications/frameworks to implement real-time analytics using open source technologies, then this book is for you. What You Will Learn Explore big data technologies and frameworks Work through practical challenges and use cases of real-time analytics versus batch analytics Develop real-word use cases for processing and analyzing data in real-time using the programming paradigm of Apache Storm Handle and process real-time transactional data Optimize and tune Apache Storm for varied workloads and production deployments Process and stream data with Amazon Kinesis and Elastic MapReduce Perform interactive and exploratory data analytics using Spark SQL Develop common enterprise architectures/applications for real-time and batch analytics In Detail Enterprise has been striving hard to deal with the challenges of data arriving in real time or near real time. Although there are technologies such

as Storm and Spark (and many more) that solve the challenges of real-time data, using the appropriate technology/framework for the right business use case is the key to success. This book provides you with the skills required to quickly design, implement and deploy your real-time analytics using real-world examples of big data use cases. From the beginning of the book, we will cover the basics of varied real-time data processing frameworks and technologies. We will discuss and explain the differences between batch and real-time processing in detail, and will also explore the techniques and programming concepts using Apache Storm. Moving on, we'll familiarize you with "Amazon Kinesis" for real-time data processing on cloud. We will further develop your understanding of real-time analytics through a comprehensive review of Apache Spark along with the high-level architecture and the building blocks of a Spark program. You will learn how to transform your data, get an output from transformations, and persist your results using Spark RDDs, using an interface called Spark SQL to work with Spark. At the end of this book, we will introduce Spark Streaming, the streaming library of Spark, and will walk you through the emerging Lambda Architecture (LA), which provides a hybrid platform for big data processing by combining real-time and precomputed batch data to provide a near real-time view of incoming data. Style and approach This step-by-step is an easy-to-follow, detailed tutorial, filled with practical examples of basic and advanced features. Each topic is explained sequentially and supported by real-world examples and executable code snippets. Microsoft Visual C# Step by Step - John Sharp

2015-10-28
Your hands-on guide to Microsoft Visual C# fundamentals with Visual Studio 2015 Expand your expertise--and teach yourself the fundamentals of programming with the latest version of Visual C# with Visual Studio 2015. If you are an experienced software developer, you'll get all the guidance, exercises, and code you need to start building responsive, scalable Windows 10 and Universal Windows Platform applications with Visual C#. Discover how to: Quickly start creating Visual C# code and projects with Visual Studio 2015 Work with variables, operators, expressions, and methods Control program flow with decision and iteration statements Build more robust apps with error, exception, and resource management Master the essentials of Visual C# object-oriented programming Use enumerations, structures, generics, collections, indexers, and other advanced features Create in-memory data queries with LINQ query expressions Improve application throughput and response time with asynchronous methods Decouple application logic and event handling Streamline development with new app templates Implement the Model-View-ViewModel (MVVM) pattern Build Universal Windows Platform apps that smoothly adapt to PCs, tablets, and Windows phones Integrate Microsoft Azure cloud databases and RESTful web services About You For software developers who are new to Visual C# or who are upgrading from older versions Readers should have experience with at least one programming language No prior Microsoft .NET or Visual Studio development experience required
Building Microservices with .NET Core 2.0 - Gaurav Aroraa 2017-12-22

Architect your .NET applications by breaking them into really small pieces - microservices -using this practical, example-based guide. Key Features Start your microservices journey and get a broader perspective on microservices development using C# 7.0 with .NET Core 2.0 Build, deploy, and test microservices using ASP.Net Core, ASP.NET Core API, and Microsoft Azure Cloud Get the basics of reactive microservices Book Description The microservices architectural style promotes the development of complex applications as a suite of small services based on business capabilities. This book will help you identify the appropriate service boundaries within your business. We'll start by looking at what microservices are and their main characteristics. Moving forward, you will be introduced to real-life application scenarios; after assessing the current issues, we will begin the journey of transforming this application by splitting it into a suite of microservices using C# 7.0 with .NET Core 2.0. You will identify service boundaries, split the application into multiple microservices, and define service contracts. You will find out how to configure, deploy, and monitor microservices, and configure scaling to allow the application to quickly adapt to increased demand in the future. With an introduction to reactive microservices, you'll strategically gain further value to keep your code base simple, focusing on what is more important rather than on messy asynchronous calls. What you will learn Get acquainted with Microsoft Azure Service Fabric Compare microservices with monolithic applications and SOA Learn Docker and Azure API management Define a service interface and implement APIs using

ASP.NET Core 2.0 Integrate services using a synchronous approach via RESTful APIs with ASP.NET Core 2.0 Implement microservices security using Azure Active Directory, OpenID Connect, and OAuth 2.0 Understand the operation and scaling of microservices in .NET Core 2.0 Understand the key features of reactive microservices and implement them using reactive extensions Who this book is for This book is for .NET Core developers who want to learn and understand the microservices architecture and implement it in their .NET Core applications. It's ideal for developers who are completely new to microservices or just have a theoretical understanding of this architectural approach and want to gain a practical perspective in order to better manage application complexities.

Designing Data-Intensive Applications - Martin Kleppmann 2017-03-16 Data is at the center of many challenges in system design today. Difficult issues need to be figured out, such as scalability, consistency, reliability, efficiency, and maintainability. In addition, we have an overwhelming variety of tools, including relational databases, NoSQL datastores, stream or batch processors, and message brokers. What are the right choices for your application? How do you make sense of all these buzzwords? In this practical and comprehensive guide, author Martin Kleppmann helps you navigate this diverse landscape by examining the pros and cons of various technologies for processing and storing data. Software keeps changing, but the fundamental principles remain the same. With this book, software engineers and architects will learn how to apply those ideas in practice, and how to make full use of data in modern applications. Peer under the

hood of the systems you already use, and learn how to use and operate them more effectively Make informed decisions by identifying the strengths and weaknesses of different tools Navigate the trade-offs around consistency, scalability, fault tolerance, and complexity Understand the distributed systems research upon which modern databases are built Peek behind the scenes of major online services, and learn from their architectures

Architectural Patterns - Pethuru Raj Chelliah 2017-12-22
Learn the importance of architectural and design patterns in producing and sustaining next-generation IT and business-critical applications with this guide. About This Book Use patterns to tackle communication, integration, application structure, and more Implement modern design patterns such as microservices to build resilient and highly available applications Choose between the MVP, MVC, and MVVM patterns depending on the application being built Who This Book Is For This book will empower and enrich IT architects (such as enterprise architects, software product architects, and solution and system architects), technical consultants, evangelists, and experts. What You Will Learn Understand how several architectural and design patterns work to systematically develop multitier web, mobile, embedded, and cloud applications Learn object-oriented and component-based software engineering principles and patterns Explore the frameworks corresponding to various architectural patterns Implement domain-driven, test-driven, and behavior-driven methodologies Deploy key platforms and tools effectively to enable EA design and solutioning Implement various patterns

designed for the cloud paradigm In Detail Enterprise Architecture (EA) is typically an aggregate of the business, application, data, and infrastructure architectures of any forward-looking enterprise. Due to constant changes and rising complexities in the business and technology landscapes, producing sophisticated architectures is on the rise. Architectural patterns are gaining a lot of attention these days. The book is divided in three modules. You'll learn about the patterns associated with object-oriented, component-based, client-server, and cloud architectures. The second module covers Enterprise Application Integration (EAI) patterns and how they are architected using various tools and patterns. You will come across patterns for Service-Oriented Architecture (SOA), Event-Driven Architecture (EDA), Resource-Oriented Architecture (ROA), big data analytics architecture, and Microservices Architecture (MSA). The final module talks about advanced topics such as Docker containers, high performance, and reliable application architectures. The key takeaways include understanding what architectures are, why they're used, and how and where architecture, design, and integration patterns are being leveraged to build better and bigger systems. Style and Approach This book adopts a hands-on approach with real-world examples and use cases. **Transactions on Large-Scale Data- and Knowledge-Centered Systems XLII** - Abdelkader Hameurlain 2019-10-17 The LNCS journal Transactions on Large-Scale Data- and Knowledge-Centered Systems focuses on data management, knowledge discovery, and knowledge processing, which are core and hot topics in computer science.

Since the 1990s, the Internet has become the main driving force behind application development in all domains. An increase in the demand for resource sharing across different sites connected through networks has led to an evolution of data- and knowledge-management systems from centralized systems to decentralized systems enabling large-scale distributed applications providing high scalability. Current decentralized systems still focus on data and knowledge as their main resource. Feasibility of these systems relies basically on P2P (peer-to-peer) techniques and the support of agent systems with scaling and decentralized control. Synergy between grids, P2P systems, and agent technologies is the key to data- and knowledge-centered systems in large-scale environments. This, the 42nd issue of Transactions on Large-Scale Data- and Knowledge-Centered Systems, consists of five revised selected regular papers, presenting the following topics: Privacy-Preserving Top-k Query Processing in Distributed Systems; Trust Factors and Insider Threats in Permissioned Distributed Ledgers: An Analytical Study and Evaluation of Popular DLT Frameworks; Polystore and Tensor Data Model for Logical Data Independence and Impedance Mismatch in Big Data Analytics; A General Framework for Multiple Choice Question Answering Based on Mutual Information and Reinforced Co-occurrence; Rejig: A Scalable Online Algorithm for Cache Server Configuration Changes. *Building Microservices Applications on Microsoft Azure* - Harsh Chawla 2019-07-17 Implement microservices starting with their architecture and moving on to their deployment, manageability, security, and monitoring. This book

focuses on the key scenarios where microservices architecture is preferred over a monolithic architecture. Building Microservices Applications on Microsoft Azure begins with a survey of microservices architecture compared to monolithic architecture and covers microservices implementation in detail. You'll see the key scenarios where microservices architecture is preferred over a monolithic approach. From there, you will explore the critical components and various deployment options of microservices on platforms such as Microsoft Azure (public cloud) and Azure Stack (hybrid cloud). This includes in-depth coverage of developing, deploying, and monitoring microservices on containers and orchestrating with Azure Service Fabric and Azure Kubernetes Cluster (AKS). This book includes practical experience from large-scale enterprise deployments, therefore it can be a quick reference for solution architects and developers to understand the critical factors while designing a microservices application. What You Will LearnExplore the use cases of microservices and monolithic architecture Discover the architecture patterns to build scalable, agile, and secure microservices applicationsDevelop and deploy microservices using Azure Service Fabric and Azure Kubernetes Service Secure microservices using the gateway patternSee the deployment options for Microservices on Azure StackImplement database patterns to handle the complexities introduced by microservices Who This Book Is For Architects and consultants who work on Microsoft Azure and manage large-scale deployments.
Cloud Native - Boris Scholl 2019-08-21
Developers often struggle when first encountering the

cloud. Learning about distributed systems, becoming familiar with technologies such as containers and functions, and knowing how to put everything together can be daunting. With this practical guide, you'll get up to speed on patterns for building cloud native applications and best practices for common tasks such as messaging, eventing, and DevOps. Authors Boris Scholl, Trent Swanson, and Peter Jausovec describe the architectural building blocks for a modern cloud native application. You'll learn how to use microservices, containers, serverless computing, storage types, portability, and functions. You'll also explore the fundamentals of cloud native applications, including how to design, develop, and operate them. Explore the technologies you need to design a cloud native application Distinguish between containers and functions, and learn when to use them Architect applications for data-related requirements Learn DevOps fundamentals and practices for developing, testing, and operating your applications Use tips, techniques, and best practices for building and managing cloud native applications Understand the costs and trade-offs necessary to make an application portable

Next Generation Databases - Guy Harrison 2015-12-30 "It's not easy to find such a generous book on big data and databases. Fortunately, this book is the one." Feng Yu. Computing Reviews. June 28, 2016. This is a book for enterprise architects, database administrators, and developers who need to understand the latest developments in database technologies. It is the book to help you choose the correct database technology at a time when concepts such as Big Data, NoSQL and NewSQL are making what used to be an easy

choice into a complex decision with significant implications. The relational database (RDBMS) model completely dominated database technology for over 20 years. Today this "one size fits all" stability has been disrupted by a relatively recent explosion of new database technologies. These paradigm-busting technologies are powering the "Big Data" and "NoSQL" revolutions, as well as forcing fundamental changes in databases across the board. Deciding to use a relational database was once truly a no-brainer, and the various commercial relational databases competed on price, performance, reliability, and ease of use rather than on fundamental architectures. Today we are faced with choices between radically different database technologies. Choosing the right database today is a complex undertaking, with serious economic and technological consequences. Next Generation Databases demystifies today's new database technologies. The book describes what each technology was designed to solve. It shows how each technology can be used to solve real word application and business problems. Most importantly, this book highlights the architectural differences between technologies that are the critical factors to consider when choosing a database platform for new and upcoming projects. Introduces the new technologies that have revolutionized the database landscape Describes how each technology can be used to solve specific application or business challenges Reviews the most popular new wave databases and how they use these new database technologies

NoSQL Data Models - Olivier Pivert 2018-07-27 The topic of NoSQL databases has recently

emerged, to face the Big Data challenge, namely the ever increasing volume of data to be handled. It is now recognized that relational databases are not appropriate in this context, implying that new database models and techniques are needed. This book presents recent research works, covering the following basic aspects: semantic data management, graph databases, and big data management in cloud environments. The chapters in this book report on research about the evolution of basic concepts such as data models, query languages, and new challenges regarding implementation issues.

NoSQL Distilled - Pramod J. Sadalage 2013
'NoSQL Distilled' is designed to provide you with enough background on how NoSQL databases work, so that you can choose the right data store without having to trawl the whole web to do it. It won't answer your questions definitively, but it should narrow down the range of options you have to consider.

Cloud Design Patterns - Alex Homer 2014-02-17
Cloud applications have a unique set of characteristics. They run on commodity hardware, provide services to untrusted users, and deal with unpredictable workloads. These factors impose a range of problems that you, as a designer or developer, need to resolve. Your applications must be resilient so that they can recover from failures, secure to protect services from malicious attacks, and elastic in order to respond to an ever changing workload. This guide demonstrates design patterns that can help you to solve the problems you might encounter in many different areas of cloud application development. Each pattern discusses design considerations, and explains how you can implement it using the features of

Windows Azure. The patterns are grouped into categories: availability, data management, design and implementation, messaging, performance and scalability, resilience, management and monitoring, and security. You will also see more general guidance related to these areas of concern. It explains key concepts such as data consistency and asynchronous messaging. In addition, there is useful guidance and explanation of the key considerations for designing features such as data partitioning, telemetry, and hosting in multiple datacenters. These patterns and guidance can help you to improve the quality of applications and services you create, and make the development process more efficient. Enjoy!

Building Microservices with .NET Core - Gaurav Kumar Aroraa 2017-06-14 Architect your .NET applications by breaking them into really small pieces—microservices—usin g this practical, example-based guide About This Book Start your microservices journey and understand a broader perspective of microservices development Build, deploy, and test microservices using ASP.Net MVC, Web API, and Microsoft Azure Cloud Get started with reactive microservices and understand the fundamentals behind it Who This Book Is For This book is for .NET Core developers who want to learn and understand microservices architecture and implement it in their .NET Core applications. It's ideal for developers who are completely new to microservices or have just a theoretical understanding of this architectural approach and want to gain a practical perspective in order to better manage application complexity. What You Will Learn Compare microservices with monolithic applications and SOA Identify the appropriate

service boundaries by mapping them to the relevant bounded contexts Define the service interface and implement the APIs using ASP.NET Web API Integrate the services via synchronous and asynchronous mechanisms Implement microservices security using Azure Active Directory, OpenID Connect, and OAuth 2.0 Understand the operations and scaling of microservices in .NET Core Understand the testing pyramid and implement consumer-driven contract using pact net core Understand what the key features of reactive microservices are and implement them using reactive extension In Detail Microservices is an architectural style that promotes the development of complex applications as a suite of small services based on business capabilities. This book will help you identify the appropriate service boundaries within the business. We'll start by looking at what microservices are, and what the main characteristics are. Moving forward, you will be introduced to real-life application scenarios, and after assessing the current issues, we will begin the journey of transforming this application by splitting it into a suite of microservices. You will identify the service boundaries, split the application into multiple microservices, and define the service contracts. You will find out how to configure, deploy, and monitor microservices, and configure scaling to allow the application to quickly adapt to increased demand in the future. With an introduction to the reactive microservices, you strategically gain further value to keep your code base simple, focusing on what is more important rather than the messy asynchronous calls. Style and approach This guide serves as a stepping stone that helps .NET Core

developers in their microservices architecture. This book provides just enough theory to understand the concepts and apply the examples. **Windows Communication Foundation 4 Step by Step** - John Sharp 2010 Your hands-on, step-by-step guide to building connected, service-oriented applications. Teach yourself the essentials of Windows Communication Foundation (WCF) 4 -- one step at a time. With this practical, learn-by-doing tutorial, you get the clear guidance and hands-on examples you need to begin creating Web services for robust Windows-based business applications. Discover how to: Build and host SOAP and REST services Maintain service contracts and data contracts Control configuration and communications programmatically Implement message encryption, authentication, and authorization Manage identity with Windows CardSpace Begin working with Windows Workflow Foundation to create scalable and durable business services Implement service discovery and message routing Optimize performance with service throttling, encoding, and streaming Integrate WCF services with ASP.NET clients and enterprise services components Your Step by Step digital content includes: Practice exercises Downloadable code samples Fully searchable online edition of the book -- with unlimited access on the Web **The Self-Service Data Roadmap** - Sandeep Uttamchandani 2020-09-10 Data-driven insights are a key competitive advantage for any industry today, but deriving insights from raw data can still take days or weeks. Most organizations can't scale data science teams fast enough to keep up with the growing amounts of data to transform. What's the answer? Self-service data.

With this practical book, data engineers, data scientists, and team managers will learn how to build a self-service data science platform that helps anyone in your organization extract insights from data. Sandeep Uttamchandani provides a scorecard to track and address bottlenecks that slow down time to insight across data discovery, transformation, processing, and production. This book bridges the gap between data scientists bottlenecked by engineering realities and data engineers unclear about ways to make self-service work. Build a self-service portal to support data discovery, quality, lineage, and governance Select the best approach for each self-service capability using open source cloud technologies Tailor self-service for the people, processes, and technology maturity of your data platform Implement capabilities to democratize data and reduce time to insight Scale your self-service portal to support a large number of users within your organization