

Databases At Scale: Operations Engineering

This volume represents a valuable collective contribution to the research and development of database systems. It contains papers in a variety of topics such as data models, distributed databases, multimedia databases, concurrency control, hypermedia and document processing, user interface, query processing and database applications. Contents: Introduction to SQL/X (W Kim)An Object-Oriented Approach to Security Policies and their Access Controls for Database Management (D K Hsiao)The ESSE Project: An Overview (R Zicari et al.)The Remote-Exchange Approach to Semantic Heterogeneity in Federated Database Systems (D McLeod)A Linear Model of Distributed Query Execution Strategies (M E Orlowska & Y-C Zhang)Multimedia Data Handling in a Knowledge Representation System (E Bertino et al.)Implementation and Evaluation of a New Approach to Storage Management for Persistent Data — Towards Virtual-Memory Databases (G-Y Bai & A Makinouchi)Hyperbase System: A Structured Architecture (R Sacks-Davis et al.)A Hypermedia Document System Based on Relational Database (S Futamura et al.)Cooperative Query Answering in CoBase (Q-M Chen & W Chu)The ADKMS Knowledge Acquisition System (E Bertino et al.)Constraints for Query Optimization in Deductive Databases (J Harland & K Ramamohanarao)The Object-Oriented Database Management — A Tutorial on its Fundamentals (D K Hsiao)and other papers Readership: Computer scientists.

Advanced Topics in Database Research is a series of books on the fields of database, software engineering, and systems analysis and design. They feature the latest research ideas and topics on how to enhance current database systems, improve information storage, refine existing database models, and develop advanced applications. Advanced Topics in Database Research, Volume 5 is a part of this series. Advanced Topics in Database Research, Volume 5 presents the latest research ideas and topics on database systems and applications, and provides insights into important developments in the field of database and database management. This book describes the capabilities and features of new technologies and methodologies, and presents state-of-the-art research ideas, with an emphasis on theoretical issues regarding databases and database management.

The infrastructure-as-code revolution in IT is also affecting database administration. With this practical book, developers, system administrators, and junior to mid-level DBAs will learn how the modern practice of site reliability engineering applies to the craft of database architecture and operations. Authors Laine Campbell and Charity Majors provide a framework for professionals looking to join the ranks of today's database reliability engineers (DBRE). You'll begin by exploring core operational concepts that DBREs need to master. Then you'll examine a wide range of database persistence options, including how to implement key technologies to provide resilient, scalable, and performant data storage and retrieval. With a firm foundation in database reliability engineering, you'll be ready to dive into the architecture and operations of any modern database. This book covers: Service-level requirements

and risk management Building and evolving an architecture for operational visibility Infrastructure engineering and infrastructure management How to facilitate the release management process Data storage, indexing, and replication Identifying datastore characteristics and best use cases Datastore architectural components and data-driven architectures

GIS and Geocomputation for Water Resource Science and Engineering not only provides a comprehensive introduction to the fundamentals of geographic information systems but also demonstrates how GIS and mathematical models can be integrated to develop spatial decision support systems to support water resources planning, management and engineering. The book uses a hands-on active learning approach to introduce fundamental concepts and numerous case-studies are provided to reinforce learning and demonstrate practical aspects. The benefits and challenges of using GIS in environmental and water resources fields are clearly tackled in this book, demonstrating how these technologies can be used to harness increasingly available digital data to develop spatially-oriented sustainable solutions. In addition to providing a strong grounding on fundamentals, the book also demonstrates how GIS can be combined with traditional physics-based and statistical models as well as information-theoretic tools like neural networks and fuzzy set theory.

The Handbook of Groundwater Engineering

Computerworld

Data Management at Scale

Seeking SRE

Future Databases '92

VLDB Workshops, Poly 2020 and DMAH 2020, Virtual Event, August 31 and September 4, 2020, Revised Selected Papers

This volume contains technical papers and panel position papers selected from the proceedings of the International Symposium on Information Systems and Technologies for Network Society, held together with the IPSJ (information processing society of Japan) National Convention, in September 1997. Papers were submitted from all over the world, especially from Japan, Korea and China. These countries are believed to form one of the major computer manufacturing centers in the world. A panel on "Computer Science Education for the 21st Century" was set up. A special session on Japanese project on Software Engineering invited representative researchers from the project. The project is supported by the Ministry of Education, Japan.

This book provides comprehensive coverage of the latest advances and trends in information technology, science and engineering. Specifically, it addresses a number of broad themes, including multi-modal informatics, data mining, agent-based and multi-agent systems for health and education informatics, which inspire the development of intelligent information technologies. The contributions cover a wide range of topics such as AI applications and innovations in health and education informatics; data and knowledge management; multi-modal application management; and web and media mining for multi-modal informatics. Outlining promising future research directions, the book is a valuable resource for students, researchers and professionals, and a useful reference guide for newcomers to the field. This book is a compilation of the papers presented in the 2021 International Conference on Multi-modal Information Analytics, held in Huhehaote, China, on April 23–24, 2021. Unrivaled coverage of a broad spectrum of industrial engineering concepts and applications The Handbook of Industrial Engineering, Third Edition contains a vast array of timely and useful methodologies for achieving increased productivity, quality, and competitiveness and improving the quality of working life in manufacturing and service industries. This astoundingly comprehensive

resource also provides a cohesive structure to the discipline of industrial engineering with four classifications: technology; performance improvement management; management, planning, and design control; and decision-making methods. Completely updated and expanded to reflect nearly a decade of important developments in the field, this Third Edition features a wealth of new information on project management, supply-chain management and logistics, and systems relationships in service industries. Other important features of this essential reference include: * More than 100 helpful tables, graphs, figures, and formulas * Step-by-step descriptions of hundreds of problem-solving methodologies * Hundreds of clear, easy-to-follow application examples * Contributions from 176 accomplished international professionals with diverse training and affiliations * More than 100 citations for further reading The Handbook of Industrial Engineering, Third Edition is an immense and useful one-stop resource for industrial engineers and technical support personnel in corporations of any size; continuous process and discrete part manufacturing industries; and all types of service industries, from healthcare to hospitality, from retailing to finance. Of related interest . . .

HANDBOOK OF HUMAN FACTORS AND ERGONOMICS, Second Edition Edited by Gavriel Salvendy (0-471-11690-4) 2,165 pages 60 chapters "A comprehensive guide that contains practical knowledge and technical background on virtually all aspects of physical, cognitive, and social ergonomics. As such, it can be a valuable source of information for any individual or organization committed to providing competitive, high-quality products and safe, productive work environments."-John F. Smith Jr., Chairman of the Board, Chief Executive Officer and President, General Motors Corporation (From the Foreword)

New Frontiers in Biomedical Engineering will be an edited work taken from the 1st Annual World Congress of Chinese Biomedical Engineers - Taipei, Taiwan 2002. As the economy develops rapidly in China and the Asian-Pacific population merges into the global healthcare system, many researchers in the West are trying to make contact with the Chinese BME scientists. At WCCBE 2002, invited leaders, materials scientists, bioengineers, molecular and cellular biologists, orthopedic surgeons, and manufacturers from P.R. of China, Taiwan, Singapore and Hong Kong covered all major BME domains: biomechanics, biomaterials and tissue engineering, medical imaging, biophotonics and instrumentation, and rehabilitation. This edited work taken from the World Congress proceedings will capture worldwide readership.

College of Engineering

Building Secure and Reliable Systems

Google Cloud Platform for Data Engineering

Hands-on Site Reliability Engineering

Library of Congress Subject Headings

The Site Reliability Workbook

For more than 20 years, Network World has been the premier provider of information, intelligence and insight for network and IT executives responsible for the digital nervous systems of large organizations. Readers are responsible for designing, implementing and managing the voice, data and video systems their companies use to support everything from business critical applications to employee collaboration and electronic commerce.

The overwhelming majority of a software system's lifespan is spent in use, not in design or implementation. So, why does conventional wisdom insist that software engineers focus primarily on the design and development of large-scale computing systems? In this collection of essays and articles, key members of Google's Site Reliability Team explain how and why their commitment to the entire lifecycle has enabled the company to successfully build, deploy, monitor, and maintain some of the largest software systems in the world. You'll learn the principles and practices that enable Google engineers to make systems more scalable, reliable, and

efficient—lessons directly applicable to your organization. This book is divided into four sections: Introduction—Learn what site reliability engineering is and why it differs from conventional IT industry practices Principles—Examine the patterns, behaviors, and areas of concern that influence the work of a site reliability engineer (SRE) Practices—Understand the theory and practice of an SRE’s day-to-day work: building and operating large distributed computing systems Management—Explore Google's best practices for training, communication, and meetings that your organization can use

Each number is the catalogue of a specific school or college of the University.

This book constitutes the refereed proceedings of the 9th International Conference on Database and Expert Systems Applications, DEXA'98, held in Vienna, Austria, in August 1998. The 81 revised full papers presented were carefully selected from a total of more than 200 submissions. The papers are organized in sections on active databases, object-oriented systems, data engineering, information retrieval, workflow and cooperative systems, spatial and temporal aspects, document management, spatial databases, adaptation and view updates, genetic algorithms, cooperative and distributed environments, interaction and communication, transaction, advanced applications, temporal aspects, oriented systems, partitioning and fragmentation, database queries, data, data warehouses, knowledge discovery and data mining, knowledge extraction, and knowledge base reduction for comprehension and reuse.

Proceedings of the World Congress for Chinese Biomedical Engineers

Building Standardized Systems Across an Engineering Organization

Application of Intelligent Systems in Multi-modal Information Analytics

Database Reliability Engineering

Practical Ways to Implement SRE

How Google Runs Production Systems

As data management and integration continue to evolve rapidly, storing all your data in one place, such as a data warehouse, is no longer scalable. In the very near future, data will need to be distributed and available for several technological solutions. With this practical book, you'll learn how to migrate your enterprise from a complex and tightly coupled data landscape to a more flexible architecture ready for the modern world of data consumption. Executives, data architects, analytics teams, and compliance and governance staff will learn how to build a modern scalable data landscape using the Scaled Architecture, which you can introduce incrementally without a large upfront investment. Author Piethein Strengholt provides blueprints, principles, observations, best practices, and patterns to get you up to speed. Examine data management trends, including technological developments, regulatory requirements, and privacy concerns Go deep into the Scaled Architecture and learn how the pieces fit together Explore data governance and data security, master data management, self-service data marketplaces, and the importance of metadata

Taking greater advantage of powerful computing capabilities over the last several years, the development of fundamental information and new models has led to major advances in nearly every aspect of chemical engineering. Albright's Chemical Engineering Handbook represents a reliable source of updated methods, applications, and fundamental concepts that will continue to play a significant role in driving new research and improving plant design and operations. Well-rounded, concise, and practical by design, this handbook collects valuable insight from an exceptional diversity of leaders in their respective specialties. Each chapter provides a clear review of basic information, case examples, and references to additional, more in-depth information. They explain essential principles, calculations, and issues relating to topics including reaction engineering, process control and design, waste disposal, and electrochemical and biochemical engineering. The final chapters cover aspects of patents and intellectual property, practical communication, and ethical considerations that are most relevant to engineers. From fundamentals to plant operations, Albright's Chemical Engineering Handbook offers a thorough, yet succinct guide to day-to-day methods and calculations used in chemical engineering applications. This handbook will serve the needs of practicing professionals as well as students preparing to enter the field.

Due to the increasing demand for adequate water supply caused by the augmenting global population, groundwater production has acquired a new importance. In many areas, surface waters are not available in sufficient quantity or quality. Thus, an increasing demand for groundwater has resulted. However, the residence time of groundwater can be of the order of thousands of years while surface waters is of the order of days. Therefore, substantially more attention is warranted for transport processes and pollution remediation in groundwater than for surface waters. Similarly, pollution remediation problems in groundwater are generally complex. This excellent, timely resource covers the field of groundwater from an engineering perspective, comprehensively addressing the range of subjects related to subsurface hydrology. It provides a practical treatment of the flow of groundwater, the transport of substances, the construction of wells and well fields, the production of groundwater, and site characterization and remediation of groundwater pollution. No other reference specializes in groundwater engineering to such a broad range of subjects. Its use extends to: The engineer designing a well or well field The engineer designing or operating a landfill facility for municipal or hazardous wastes The hydrogeologist

investigating a contaminant plume The engineer examining the remediation of a groundwater pollution problem The engineer or lawyer studying the laws and regulations related to groundwater quality The scientist analyzing the mechanics of solute transport The geohydrologist assessing the regional modeling of aquifers The geophysicist determining the characterization of an aquifer The cartographer mapping aquifer characteristics The practitioner planning a monitoring network

A guide to the development and manufacturing of pharmaceutical products written for professionals in the industry, revised second edition The revised and updated second edition of Chemical Engineering in the Pharmaceutical Industry is a practical book that highlights chemistry and chemical engineering. The book's regulatory quality strategies target the development and manufacturing of pharmaceutically active ingredients of pharmaceutical products. The expanded second edition contains revised content with many new case studies and additional example calculations that are of interest to chemical engineers. The 2nd Edition is divided into two separate books: 1) Active Pharmaceutical Ingredients (API's) and 2) Drug Product Design, Development and Modeling. The active pharmaceutical ingredients book puts the focus on the chemistry, chemical engineering, and unit operations specific to development and manufacturing of the active ingredients of the pharmaceutical product. The drug substance operations section includes information on chemical reactions, mixing, distillations, extractions, crystallizations, filtration, drying, and wet and dry milling. In addition, the book includes many applications of process modeling and modern software tools that are geared toward batch-scale and continuous drug substance pharmaceutical operations. This updated second edition:

- Contains 30 new chapters or revised chapters specific to API, covering topics including: manufacturing quality by design, computational approaches, continuous manufacturing, crystallization and final form, process safety
- Expanded topics of scale-up, continuous processing, applications of thermodynamics and thermodynamic modeling, filtration and drying
- Presents updated and expanded example calculations
- Includes contributions from noted experts in the field

Written for pharmaceutical engineers, chemical engineers, undergraduate and graduate students, and professionals in the field of pharmaceutical sciences and manufacturing, the second edition of Chemical Engineering in the Pharmaceutical Industry focuses on the development and chemical engineering as well as operations specific to the design, formulation, and manufacture of drug substance and products.

Database and Expert Systems Applications

Handbook of Industrial Engineering

Chemical Engineering in the Pharmaceutical Industry, Active

Pharmaceutical Ingredients

Site Reliability Engineering

GIS and Geocomputation for Water Resource Science and
Engineering

Principles and Modeling

Information Organization and Databases: Foundations of Data

Organization provides recent developments of information organization technologies that have become crucial not only for data mining applications and information visualization, but also for treatment of semistructured data, spatio-temporal data and multimedia data that are not necessarily stored in conventional DBMSs. Information Organization and Databases: Foundations of Data Organization presents: semistructured data addressing XML, query languages and integrity constraints, focusing on advanced technologies for organizing web data for effective retrieval; multimedia database organization emphasizing video data organization and data structures for similarity retrieval; technologies for data mining and data warehousing; index organization and efficient query processing issues; spatial data access and indexing; organizing and retrieval of WWW and hypermedia. Information Organization and Databases: Foundations of Data Organization is a resource for database practitioners, database researchers, designers and administrators of multimedia information systems, and graduate-level students in the area of information retrieval and/or databases wishing to keep abreast of advances in the information organization technologies.

In 2016, Google's Site Reliability Engineering book ignited an industry discussion on what it means to run production services today—and why reliability considerations are fundamental to service design. Now, Google engineers who worked on that bestseller introduce The Site Reliability Workbook, a hands-on companion that uses concrete examples to show you how to put SRE principles and practices to work in your environment. This new workbook not only combines practical examples from Google's experiences, but also provides case studies from Google's Cloud Platform customers who underwent this journey. Evernote, The Home Depot, The New York Times, and other companies outline hard-won experiences of what worked for them and what didn't. Dive into this workbook and learn how to flesh out your own SRE practice, no matter what size your company is. You'll learn: How to run reliable services in environments you don't completely control—like cloud Practical applications of how to create, monitor, and run your services via Service Level Objectives How to convert existing ops teams to SRE—including how to dig out of operational overload Methods for starting SRE from either greenfield or

brownfield

Can a system be considered truly reliable if it isn't fundamentally secure? Or can it be considered secure if it's unreliable? Security is crucial to the design and operation of scalable systems in production, as it plays an important part in product quality, performance, and availability. In this book, experts from Google share best practices to help your organization design scalable and reliable systems that are fundamentally secure. Two previous O'Reilly books from Google—Site Reliability Engineering and The Site Reliability Workbook—demonstrated how and why a commitment to the entire service lifecycle enables organizations to successfully build, deploy, monitor, and maintain software systems. In this latest guide, the authors offer insights into system design, implementation, and maintenance from practitioners who specialize in security and reliability. They also discuss how building and adopting their recommended best practices requires a culture that's supportive of such change. You'll learn about secure and reliable systems through:

- Design strategies**
- Recommendations for coding, testing, and debugging practices**
- Strategies to prepare for, respond to, and recover from incidents**
- Cultural best practices that help teams across your organization collaborate effectively**

Today, software engineers need to know not only how to program effectively but also how to develop proper engineering practices to make their codebase sustainable and healthy. This book emphasizes this difference between programming and software engineering. How can software engineers manage a living codebase that evolves and responds to changing requirements and demands over the length of its life? Based on their experience at Google, software engineers Titus Winters and Hyrum Wright, along with technical writer Tom Manshreck, present a candid and insightful look at how some of the world's leading practitioners construct and maintain software. This book covers Google's unique engineering culture, processes, and tools and how these aspects contribute to the effectiveness of an engineering organization. You'll explore three fundamental principles that software organizations should keep in mind when designing, architecting, writing, and maintaining code:

- How time affects the sustainability of software and how to make your code resilient over time**
- How scale affects the viability of software practices within an engineering organization**
- What trade-offs a typical engineer needs to make when evaluating design and development decisions**

Environmental Process Analysis

Information Organization and Databases

Heterogeneous Data Management, Polystores, and Analytics for Healthcare

Best Practices for Designing, Implementing, and Maintaining Systems

Framework for a National Database System for Maintenance Actions on

Highway Bridges

Build Capability to Design, Deploy, Monitor, and Sustain Enterprise Software Systems at Scale (English Edition)

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

This book constitutes revised selected papers from two VLDB workshops: The International Workshop on Polystore Systems for Heterogeneous Data in Multiple Databases with Privacy and Security Assurances, Poly 2020, and the 6th International Workshop on Data Management and Analytics for Medicine and Healthcare, DMAH 2020, which were held virtually on August 31 and September 4, 2020. For Poly 2020, 4 full and 3 short papers were accepted from 10 submissions; and for DMAH 2020, 7 full and 2 short papers were accepted from a total of 15 submissions. The papers were organized in topical sections as follows: Privacy, Security and/or Policy Issues for Heterogeneous Data; COVID-19 Data Analytics and Visualization; Deep Learning based Biomedical Data Analytics; NLP based Learning from Unstructured Data; Biomedical Data Modelling and Prediction.

Organizations big and small have started to realize just how crucial system and application reliability is to their business. They've also learned just how difficult it is to maintain that reliability while iterating at the speed demanded by the marketplace. Site Reliability Engineering (SRE) is a proven approach to this challenge. SRE is a large and rich topic to discuss. Google led the way with Site Reliability Engineering, the wildly successful O'Reilly book that described Google's creation of the discipline and the implementation that's allowed them to operate at a planetary scale. Inspired by that earlier work, this book explores a very different part of the SRE space. The more than two dozen chapters in Seeking SRE bring you into some of the important conversations going on in the SRE world right now. Listen as engineers and other leaders in the field discuss: Different ways of implementing SRE and SRE principles in a wide variety of settings How SRE relates to other approaches such as DevOps Specialties on the cutting edge that will soon be commonplace in SRE Best practices and technologies that make practicing SRE easier The important but rarely explored human side of SRE David N. Blank-Edelman is the book's curator and editor.

Google Cloud Platform for Data Engineering is designed to take the beginner through a journey to become a competent and certified GCP data engineer. The book, therefore, is split into three parts; the first part covers fundamental concepts of data engineering and data analysis from a platform and technology-neutral perspective. Reading part 1 will bring a beginner up to speed

with the generic concepts, terms and technologies we use in data engineering. The second part, which is a high-level but comprehensive introduction to all the concepts, components, tools and services available to us within the Google Cloud Platform. Completing this section will provide the beginner to GCP and data engineering with a solid foundation on the architecture and capabilities of the GCP. Part 3, however, is where we delve into the moderate to advanced techniques that data engineers need to know and be able to carry out. By this time the raw beginner you started the journey at the beginning of part 1 will be a knowledgeable albeit inexperienced data engineer. However, by the conclusion of part 3, they will have gained the advanced knowledge of data engineering techniques and practices on the GCP to pass not only the certification exam but also most interviews and practical tests with confidence. In short part 3, will provide the prospective data engineer with detailed knowledge on setting up and configuring DataProc - GCP's version of the Spark/Hadoop ecosystem for big data. They will also learn how to build and test streaming and batch data pipelines using pub/sub/ dataFlow and BigQuery. Furthermore, they will learn how to integrate all the ML and AI Platform components and APIs. They will be accomplished in connecting data analysis and visualisation tools such as Datalab, DataStudio and AI notebooks amongst others. They will also by now know how to build and train a TensorFlow DNN using APIs and Keras and optimise it to run large public data sets. Also, they will know how to provision and use Kubeflow and Kube Pipelines within Google Kubernetes engines to run container workloads as well as how to take advantage of serverless technologies such as Cloud Run and Cloud Functions to build transparent and seamless data processing platforms. The best part of the book though is its compartmental design which means that anyone from a beginner to an intermediate can join the book at whatever point they feel comfortable.

Production-Ready Microservices

Learning Spark

From Beginner to Data Engineer using Google Cloud Platform

The Big Ideas Behind Reliable, Scalable, and Maintainable Systems

The Chemical Engineer

InfoWorld

One of the biggest challenges for organizations that have adopted microservice architecture is the lack of architectural, operational, and organizational standardization. After splitting a monolithic application or building a microservice ecosystem from scratch, many engineers are left wondering what's next. In this practical book, author Susan Fowler presents a set of microservice standards in depth, drawing from her experience standardizing over a thousand microservices at Uber. You'll learn how to design microservices that are stable, reliable, scalable, fault tolerant, performant, monitored, documented, and prepared for any catastrophe. Explore production-readiness standards, including: Stability and Reliability: develop, deploy, introduce, and deprecate microservices; protect against dependency failures Scalability and Performance: learn essential components for achieving

greater microservice efficiency Fault Tolerance and Catastrophe Preparedness: ensure availability by actively pushing microservices to fail in real time Monitoring: learn how to monitor, log, and display key metrics; establish alerting and on-call procedures Documentation and Understanding: mitigate tradeoffs that come with microservice adoption, including organizational sprawl and technical debt

Although service-level objectives (SLOs) continue to grow in importance, there's a distinct lack of information about how to implement them. Practical advice that does exist usually assumes that your team already has the infrastructure, tooling, and culture in place. In this book, recognized SLO expert Alex Hidalgo explains how to build an SLO culture from the ground up. Ideal as a primer and daily reference for anyone creating both the culture and tooling necessary for SLO-based approaches to reliability, this guide provides detailed analysis of advanced SLO and service-level indicator (SLI) techniques. Armed with mathematical models and statistical knowledge to help you get the most out of an SLO-based approach, you'll learn how to build systems capable of measuring meaningful SLIs with buy-in across all departments of your organization. Define SLIs that meaningfully measure the reliability of a service from a user's perspective Choose appropriate SLO targets, including how to perform statistical and probabilistic analysis Use error budgets to help your team have better discussions and make better data-driven decisions Build supportive tooling and resources required for an SLO-based approach Use SLO data to present meaningful reports to leadership and your users

Data is bigger, arrives faster, and comes in a variety of formats—and it all needs to be processed at scale for analytics or machine learning. But how can you process such varied workloads efficiently? Enter Apache Spark. Updated to include Spark 3.0, this second edition shows data engineers and data scientists why structure and unification in Spark matters. Specifically, this book explains how to perform simple and complex data analytics and employ machine learning algorithms. Through step-by-step walk-throughs, code snippets, and notebooks, you'll be able to: Learn Python, SQL, Scala, or Java high-level Structured APIs Understand Spark operations and SQL Engine Inspect, tune,

and debug Spark operations with Spark configurations and Spark UI Connect to data sources: JSON, Parquet, CSV, Avro, ORC, Hive, S3, or Kafka Perform analytics on batch and streaming data using Structured Streaming Build reliable data pipelines with open source Delta Lake and Spark Develop machine learning pipelines with MLlib and productionize models using MLflow

TRB's National Cooperative Highway Research Program (NCHRP) Report 668: Framework for a National Database System for Maintenance Actions on Highway Bridges explores a potential framework that provides a uniform format for collecting, reporting, and storing information on bridge maintenance actions for inclusion in a national bridge maintenance database. Appendixes A through E to NCHRP Report 668 provide detailed information on the different aspects of the research. Appendix A: Information on Bridge Maintenance Programs; Appendix B: National Bridge Maintenance Database Tables; Appendix C: List of Element Level Costs of Maintenance Actions; Appendix D: Examples of National Bridge Maintenance Database Uses; Appendix E: Other National Bridge Maintenance Database Tables--

Network World

Implementing Service Level Objectives

9th International Conference, DEXA'98, Vienna, Austria, August 24-28, 1998, Proceedings

Lessons Learned from Programming Over Time

Technology and Operations Management

Combines language tutorials with application design advice to cover the PHP server-side scripting language and the MySQL database engine.

Enables readers to apply core principles of environmental engineering to analyze environmental systems Environmental Process Analysis takes a unique approach, applying mathematical and numerical process modeling within the context of both natural and engineered environmental systems. Readers master core principles of natural and engineering science such as chemical equilibria, reaction kinetics, ideal and non-ideal reactor theory, and mass accounting by performing practical real-world analyses. As they progress through the text, readers will have the opportunity to analyze a broad range of environmental processes and systems, including water and wastewater treatment, surface mining, agriculture, landfills, subsurface saturated and unsaturated porous media, aqueous and marine sediments, surface waters, and atmospheric moisture. The text begins with an examination of water, core definitions, and a review of important chemical principles. It then progressively builds upon this base with applications of Henry's law, acid/base equilibria, and reactions in ideal reactors. Finally, the text addresses reactions in non-ideal

reactors and advanced applications of acid/base equilibria, complexation and solubility/dissolution equilibria, and oxidation/reduction equilibria. Several tools are provided to fully engage readers in mastering new concepts and then applying them in practice, including: Detailed examples that demonstrate the application of concepts and principles Problems at the end of each chapter challenging readers to apply their newfound knowledge to analyze environmental processes and systems MathCAD worksheets that provide a powerful platform for constructing process models Environmental Process Analysis serves as a bridge between introductory environmental engineering textbooks and hands-on environmental engineering practice. By learning how to mathematically and numerically model environmental processes and systems, readers will also come to better understand the underlying connections among the various models, concepts, and systems.

The infrastructure-as-code revolution in IT is also affecting database administration. With this practical book, developers, system administrators, and junior to mid-level DBAs will learn how the modern practice of site reliability engineering applies to the craft of database architecture and operations. Authors Laine Campbell and Charity Majors provide a framework for professionals looking to join the ranks of today's database reliability engineers (DBRE). You'll begin by exploring core operational concepts that DBREs need to master. Then you'll examine a wide range of database persistence options, including how to implement key technologies to provide resilient, scalable, and performant data storage and retrieval. With a firm foundation in database reliability engineering, you'll be ready to dive into the architecture and operations of any modern database. This book covers: Service-level requirements and risk management Building and evolving an architecture for operational visibility Infrastructure engineering and infrastructure management How to facilitate the release management process Data storage, indexing, and replication Identifying datastore characteristics and best use cases Datastore architectural components and data-driven architectures

For more than 40 years, Computerworld has been the leading source of technology news and information for IT influencers worldwide. Computerworld's award-winning Web site (Computerworld.com), twice-monthly publication, focused conference series and custom research form the hub of the world's largest global IT media network.

Designing Data-Intensive Applications

Conversations About Running Production Systems at Scale

Designing and Operating Resilient Database Systems

2021 International Conference on Multi-modal Information Analytics (MMIA 2021), Volume 1

Foundations of Data Organization

Web Database Applications with PHP and MySQL

Database Reliability Engineering***Designing and Operating Resilient Database Systems***

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

A comprehensive guide with basic to advanced SRE practices and hands-on

examples. KEY FEATURES ? Demonstrates how to execute site reliability engineering along with fundamental concepts. ? Illustrates real-world examples and successful techniques to put SRE into production. ? Introduces you to DevOps, advanced techniques of SRE, and popular tools in use. DESCRIPTION Hands-on Site Reliability Engineering (SRE) brings you a tailor-made guide to learn and practice the essential activities for the smooth functioning of enterprise systems, right from designing to the deployment of enterprise software programs and extending to scalable use with complete efficiency and reliability. The book explores the fundamentals around SRE and related terms, concepts, and techniques that are used by SRE teams and experts. It discusses the essential elements of an IT system, including microservices, application architectures, types of software deployment, and concepts like load balancing. It explains the best techniques in delivering timely software releases using containerization and CI/CD pipeline. This book covers how to track and monitor application performance using Grafana, Prometheus, and Kibana along with how to extend monitoring more effectively by building full-stack observability into the system. The book also talks about chaos engineering, types of system failures, design for high-availability, DevSecOps and AIOps. WHAT YOU WILL LEARN ? Learn the best techniques and practices for building and running reliable software. ? Explore observability and popular methods for effective monitoring of applications. ? Workaround SLIs, SLOs, Error Budgets, and Error Budget Policies to manage failures. ? Learn to practice continuous software delivery using blue/green and canary deployments. ? Explore chaos engineering, SRE best practices, DevSecOps and AIOps. WHO THIS BOOK IS FOR This book caters to experienced IT professionals, application developers, software engineers, and all those who are looking to develop SRE capabilities at the individual or team level. TABLE OF CONTENTS 1. Understand the World of IT 2. Introduction to DevOps 3. Introduction to SRE 4. Identify and Eliminate Toil 5. Release Engineering 6. Incident Management 7. IT Monitoring 8. Observability 9. Key SRE KPIs: SLAs, SLOs, SLIs, and Error Budgets 10. Chaos Engineering 11. DevSecOps and AIOps 12. Culture of Site Reliability Engineering

University of Michigan Official Publication
Albright's Chemical Engineering Handbook
Software Engineering at Google
Frontiers in Biomedical Engineering
Information Systems And Technologies For Network Society: Proceedings Of The Ipsj International Symposium
Advanced Topics in Database Research, Volume 5