

Time Series Databases New Ways To Store And Access Data

Thank you very much for downloading **Time Series Databases New Ways To Store And Access Data** . Maybe you have knowledge that, people have look numerous times for their favorite books like this Time Series Databases New Ways To Store And Access Data , but end up in harmful downloads. Rather than reading a good book with a cup of coffee in the afternoon, instead they are facing with some infectious virus inside their computer.

Time Series Databases New Ways To Store And Access Data is available in our digital library an online access to it is set as public so you can download it instantly.

Our digital library hosts in multiple countries, allowing you to get the most less latency time to download any of our books like this one.

Merely said, the Time Series Databases New Ways To Store And Access Data is universally compatible with any devices to read

Real-World Hadoop - Ted Dunning 2015-03-24
If you're a business team leader, CIO, business analyst, or developer interested in how Apache Hadoop and Apache HBase-related technologies can

address problems involving large-scale data in cost-effective ways, this book is for you. Using real-world stories and situations, authors Ted Dunning and Ellen Friedman show Hadoop newcomers and

seasoned users alike how NoSQL databases and Hadoop can solve a variety of business and research issues. You'll learn about early decisions and pre-planning that can make the process easier and more productive. If you're already using these technologies, you'll discover ways to gain the full range of benefits possible with Hadoop. While you don't need a deep technical background to get started, this book does provide expert guidance to help managers, architects, and practitioners succeed with their Hadoop projects. Examine a day in the life of big data: India's ambitious Aadhaar project Review tools in the Hadoop ecosystem such as Apache's Spark, Storm, and Drill to learn how they can help you Pick up a collection of technical and strategic tips that have helped others succeed with Hadoop Learn from several prototypical Hadoop use cases, based on how organizations have actually applied the technology Explore real-world stories that reveal how MapR customers

combine use cases when putting Hadoop and NoSQL to work, including in production *Programming Hive* - Edward Capriolo 2012-09-26 Describes the features and functions of Apache Hive, the data infrastructure for Hadoop. *Environmental Software Systems. Infrastructures, Services and Applications* - Ralf Denzer 2015-02-09 This book constitutes the refereed proceedings of the 11th IFIP WG 5.11 International Symposium on Environmental Software Systems, ISESS 2015, held in Melbourne, Australia, in March 2015. The 62 revised full papers presented were carefully reviewed and selected from 104 submissions. The papers are organized in the following topical sections: information systems, information modeling and semantics; decision support tools and systems; modelling and simulation systems; architectures, infrastructures, platforms and services; requirements, software engineering and software tools;

analytics and visualization; and high-performance computing and big data.

Principles of Database Management - Wilfried Lemahieu 2018-07-12

Introductory, theory-practice balanced text teaching the fundamentals of databases to advanced undergraduates or graduate students in information systems or computer science.

Energy Time Series Forecasting - Lars Dannecker 2015-08-06

Lars Dannecker developed a novel online forecasting process that significantly improves how forecasts are calculated. It increases forecasting efficiency and accuracy, as well as allowing the process to adapt to different situations and applications. Improving the forecasting efficiency is a key pre-requisite for ensuring stable electricity grids in the face of an increasing amount of renewable energy sources. It is also important to facilitate the move from static day ahead electricity trading towards

more dynamic real-time marketplaces. The online forecasting process is realized by a number of approaches on the logical as well as on the physical layer that we introduce in the course of this book. Nominated for the Georg-Helm-Preis 2015 awarded by the Technische Universität Dresden.

Storage Systems - Alexander Thomasian 2021-10-13
Storage Systems: Organization, Performance, Coding, Reliability and Their Data Processing was motivated by the 1988 Redundant Array of Inexpensive/Independent Disks proposal to replace large form factor mainframe disks with an array of commodity disks. Disk loads are balanced by striping data into strips—with one strip per disk—and storage reliability is enhanced via replication or erasure coding, which at best dedicates k strips per stripe to tolerate k disk failures. Flash memories have resulted in a paradigm shift with Solid State Drives (SSDs) replacing Hard Disk Drives (HDDs) for high performance

applications. RAID and Flash have resulted in the emergence of new storage companies, namely EMC, NetApp, SanDisk, and Purestorage, and a multibillion-dollar storage market. Key new conferences and publications are reviewed in this book. The goal of the book is to expose students, researchers, and IT professionals to the more important developments in storage systems, while covering the evolution of storage technologies, traditional and novel databases, and novel sources of data. We describe several prototypes: FAWN at CMU, RAMCloud at Stanford, and Lightstore at MIT; Oracle's Exadata, AWS' Aurora, Alibaba's PolarDB, Fungible Data Center; and author's paper designs for cloud storage, namely heterogeneous disk arrays and hierarchical RAID. • Surveys storage technologies and lists sources of data: measurements, text, audio, images, and video • Familiarizes with paradigms to improve performance: caching,

prefetching, log-structured file systems, and merge-trees (LSMs) • Describes RAID organizations and analyzes their performance and reliability • Conserves storage via data compression, deduplication, compaction, and secures data via encryption • Specifies implications of storage technologies on performance and power consumption • Exemplifies database parallelism for big data, analytics, deep learning via multicore CPUs, GPUs, FPGAs, and ASICs, e.g., Google's Tensor Processing Units

Monitoring with Graphite -
Jason Dixon 2017-03-13

Graphite has become one of the most powerful monitoring tools available today, due to its ease of use, rapid graph prototyping abilities, and a friendly rendering API. With this practical guide, system administrators and engineers will learn how to use this open source tool to track operational data you need to monitor your systems, as well as application-level metrics for profiling your

services. Author Jason Dixon, member of the Graphite project, provides a thorough introduction of Graphite from the basics to the skills and tools you need for troubleshooting and scaling out its software components. If you want to learn more about monitoring systems, services, or applications, this is the book you need. Get an introduction to monitoring, including important concepts and terminology Examine the features and functionality of key Graphite components, including Carbon and Whisper Learn the typical user workflow necessary to create a basic line chart Build complex charts with chained functions and multiple axes that interact directly with the rendering API Understand how to use the native Graphite dashboard, as well as the more popular third-party dashboards Master the art of scaling and troubleshooting high-performance or highly available Graphite clusters

The Cloud-Based Demand-Driven Supply Chain - Vinit

Sharma 2018-11-06

It's time to get your head in the cloud! In today's business environment, more and more people are requesting cloud-based solutions to help solve their business challenges. So how can you not only anticipate your clients' needs but also keep ahead of the curve to ensure their goals stay on track? With the help of this accessible book, you'll get a clear sense of cloud computing and understand how to communicate the benefits, drawbacks, and options to your clients so they can make the best choices for their unique needs. Plus, case studies give you the opportunity to relate real-life examples of how the latest technologies are giving organizations worldwide the opportunity to thrive as supply chain solutions in the cloud. Demonstrates how improvements in forecasting, collaboration, and inventory optimization can lead to cost savings Explores why cloud computing is becoming increasingly important Takes a close look at the types of cloud

computing Makes sense of demand-driven forecasting using Amazon's cloud Whether you work in management, business, or IT, this is the dog-eared reference you'll want to keep close by as you continue making sense of the cloud.

Theory and Applications of Time Series Analysis - Olga Valenzuela 2020-11-20

This book presents a selection of peer-reviewed contributions on the latest advances in time series analysis, presented at the International Conference on Time Series and Forecasting (ITISE 2019), held in Granada, Spain, on September 25-27, 2019. The first two parts of the book present theoretical contributions on statistical and advanced mathematical methods, and on econometric models, financial forecasting and risk analysis. The remaining four parts include practical contributions on time series analysis in energy; complex/big data time series and forecasting; time series analysis with computational intelligence; and time series analysis and prediction for

other real-world problems. Given this mix of topics, readers will acquire a more comprehensive perspective on the field of time series analysis and forecasting. The ITISE conference series provides a forum for scientists, engineers, educators and students to discuss the latest advances and implementations in the foundations, theory, models and applications of time series analysis and forecasting. It focuses on interdisciplinary research encompassing computer science, mathematics, statistics and econometrics.

A Comparison of NoSQL Time Series Databases -

Kevin Rudolph 2015-05-21
Research Paper
(undergraduate) from the year 2015 in the subject Engineering - Industrial Engineering and Management, grade: 1,0, Technical University of Berlin (Wirtschaftsinformatik - Information Systems Engineering (ISE)), course: Seminar: Hot Topics in Information Systems

Engineering, language: English, abstract: During the last years NoSQL databases have been developed to address the needs of tremendous performance, reliability and horizontal scalability. NoSQL time series databases (TSDBs) have risen to combine valuable NoSQL properties with characteristics of time series data encountering many use-cases. Solutions offer the efficient handling of data volume and frequency related to time series. Developers and decision makers struggle with the choice of a TSDB among a large variety of solutions. Up to now no comparison exists focusing on the specific features and qualities of those heterogeneous applications. This paper aims to deliver two frameworks for the comparison of TSDBs, firstly with a focus on features and secondly on quality. Furthermore, we apply and evaluate the frameworks on up to seven open-source TSDBs such as InfluxDB and OpenTSDB. We come to the result that the investigated TSDBs differ mainly in support-

and extension related points. They share performance-enhancing techniques, time-related query capabilities and data schemas optimized for the handling of time-series data.

Neural Information

Processing - Haiqin Yang
2020-11-18

The three-volume set of LNCS 12532, 12533, and 12534 constitutes the proceedings of the 27th International Conference on Neural Information Processing, ICONIP 2020, held in Bangkok, Thailand, in November 2020. Due to COVID-19 pandemic the conference was held virtually. The 187 full papers presented were carefully reviewed and selected from 618 submissions. The papers address the emerging topics of theoretical research, empirical studies, and applications of neural information processing techniques across different domains. The third volume, LNCS 12534, is organized in topical sections on biomedical information; neural data analysis; neural network models; recommender systems;

time series analysis.

Real-Time & Stream Data

Management - Wolfram

Wingerath 2019-01-02

While traditional databases excel at complex queries over historical data, they are inherently pull-based and therefore ill-equipped to push new information to clients. Systems for data stream management and processing, on the other hand, are natively push-oriented and thus facilitate reactive behavior. However, they do not retain data indefinitely and are therefore not able to answer historical queries. The book provides an overview over the different (push-based) mechanisms for data retrieval in each system class and the semantic differences between them. It also provides a comprehensive overview over the current state of the art in real-time databases. It first includes an in-depth system survey of today's real-time databases: Firebase, Meteor, RethinkDB, Parse, Baqend, and others. Second, the high-level classification scheme

illustrated above provides a gentle introduction into the

system space of data management: Abstracting from the extreme system diversity in this field, it helps readers build a mental model of the available options.

Practical Time Series

Analysis - Aileen Nielsen

2019-09-20

Time series data analysis is increasingly important due to the massive production of such data through the internet of things, the digitalization of healthcare, and the rise of smart cities. As continuous monitoring and data collection become more common, the need for competent time series analysis with both statistical and machine learning techniques will increase. Covering innovations in time series data analysis and use cases from the real world, this practical guide will help you solve the most common data engineering and analysis challenges in time series, using both traditional statistical and modern machine learning techniques. Author Aileen

Nielsen offers an accessible, well-rounded introduction to time series in both R and Python that will have data scientists, software engineers, and researchers up and running quickly. You'll get the guidance you need to confidently: Find and wrangle time series data Undertake exploratory time series data analysis Store temporal data Simulate time series data Generate and select features for a time series Measure error Forecast and classify time series with machine or deep learning Evaluate accuracy and performance

The DevOps 2.1 Toolkit:

Docker Swarm - Viktor Farcic
2017-05-10

Viktor Farcic's latest book, The DevOps 2.1 Toolkit: Docker Swarm, shows you how to successfully integrate Docker Swarm into your DevOps toolset. About This Book Expand your DevOps Toolkit with the DevOps thought leader, Viktor Farcic Build, test, deploy, and monitor services inside Docker Swarm clusters Translate your

understanding to different hosting providers like AWS, Azure, and DigitalOcean Go beyond simple deployment to explore how to create a continuous deployment process Extend the deep understanding you gained from Viktor's DevOps 2.0 Toolkit book Who This Book Is For This book is for professionals interested in the full microservices life cycle combined with continuous deployment and containers. Target audience could be architects who want to know how to design their systems around microservices. It could be DevOps wanting to know how to apply modern configuration management practices and continuously deploy applications packed in containers. It is for developers who would like to take the process back into their hands as well as for managers who would like to gain a better understanding of the process used to deliver software from the beginning to the end. This book is for everyone wanting to know more about the software development life cycle starting

from requirements and design, through the development and testing all the way until deployment and post-deployment phases. We'll create the processes taking into account the best practices developed by and for some of the biggest companies. What You Will Learn Learn all aspects of Docker Swarm from building, testing, deploying, and monitoring services inside Docker Swarm clusters, available since Docker 1.12. Master the deeper logic of DevOps with Viktor, so that you can successfully apply that logic across any specific set of tools you're working with. Translate a deep understanding to different hosting providers like AWS, Azure, DigitalOcean, among others. You'll go beyond simple deployment: you will explore with Viktor how to create a continuous deployment process. Accomplish zero-downtime deployments, and what to do in case of a failover. Know how to run services at scale, how to monitor the systems, and how to make it

heal itself. In Detail Viktor Farcic's latest book, The DevOps 2.1 Toolkit: Docker Swarm, takes you deeper into one of the major subjects of his international best seller, The DevOps 2.0 Toolkit, and shows you how to successfully integrate Docker Swarm into your DevOps toolset. Viktor shares with you his expert knowledge in all aspects of building, testing, deploying, and monitoring services inside Docker Swarm clusters. You'll go through all the tools required for running a cluster. You'll travel through the whole process with clusters running locally on a laptop. Once you're confident with that outcome, Viktor shows you how to translate your experience to different hosting providers like AWS, Azure, and DigitalOcean. Viktor has updated his DevOps 2.0 framework in this book to use the latest and greatest features and techniques introduced in Docker. We'll go through many practices and even more tools. While there will be a lot of theory, this is a hands-on book. You won't be

able to complete it by reading it on the metro on your way to work. You'll have to read this book while in front of the computer and get your hands dirty. Style and approach We'll go through many practices and even more tools. While there will be a lot of theory, this is a hands-on book. You'll have to read this book while in front of the computer and get your hands dirty. The goal is not to master one particular set of tools, but to learn the logic behind them so that you can apply it to your job in various contexts.

Security, Privacy, and Anonymity in Computation, Communication, and Storage -
Guojun Wang 2017-12-11

This book constitutes the refereed proceedings of 11 symposia and workshops held at the 10th International Conference on Security, Privacy and Anonymity in Computation, Communication, and Storage, SpaCCS 2017, held in Guangzhou, China, in December 2017. The total of 75 papers presented in this volume was carefully reviewed

and selected from a total of 190 submissions to all workshops: UbiSafe 2017: The 9th IEEE International Symposium on UbiSafe Computing ISSR 2017: The 9th IEEE International Workshop on Security in e-Science and e-Research TrustData 2017: The 8th International Workshop on Trust, Security and Privacy for Big Data TSP 2017: The 7th International Symposium on Trust, Security and Privacy for Emerging Applications SPIoT 2017: The 6th International Symposium on Security and Privacy on Internet of Things NOPE 2017: The 5th International Workshop on Network Optimization and Performance Evaluation DependSys 2017: The Third International Symposium on Dependability in Sensor, Cloud, and Big Data Systems and Applications SCS 2017: The Third International Symposium on Sensor-Cloud Systems WCSSC 2017: The Second International Workshop on Cloud Storage Service and Computing MSCF 2017: The First International Symposium

on Multimedia Security and Digital Forensics SPBD 2017: The 2017 International Symposium on Big Data and Machine Learning in Information Security, Privacy and Anonymity

Performance Evaluation and Benchmarking - Raghunath Nambiar 2022

This book constitutes the refereed post-conference proceedings of the 13th TPC Technology Conference on Performance Evaluation and Benchmarking, TPCTC 2021, held in August 2021. The 9 papers presented were carefully reviewed and selected from numerous submissions. The TPC encourages researchers and industry experts to present and debate novel ideas and methodologies in performance evaluation, measurement, and characterization.

Time Series Databases: New Ways to Store and Access Data - Ted Dunning
2014-11-28

Time series data is of growing importance, especially with the rapid expansion of the Internet

of Things. This concise guide shows you effective ways to collect, persist, and access large-scale time series data for analysis. You'll explore the theory behind time series databases and learn practical methods for implementing them. Authors Ted Dunning and Ellen Friedman provide a detailed examination of open source tools such as OpenTSDB and new modifications that greatly speed up data ingestion. You'll learn: A variety of time series use cases The advantages of NoSQL databases for large-scale time series data NoSQL table design for high-performance time series databases The benefits and limitations of OpenTSDB How to access data in OpenTSDB using R, Go, and Ruby How time series databases contribute to practical machine learning projects How to handle the added complexity of geo-temporal data For advice on analyzing time series data, check out Practical Machine Learning: A New Look at Anomaly Detection, also from

Ted Dunning and Ellen Friedman.

Euro-Par 2020: Parallel Processing - Maciej Malawski
2020-08-18

This book constitutes the proceedings of the 26th International Conference on Parallel and Distributed Computing, Euro-Par 2020, held in Warsaw, Poland, in August 2020. The conference was held virtually due to the coronavirus pandemic. The 39 full papers presented in this volume were carefully reviewed and selected from 158 submissions. They deal with parallel and distributed computing in general, focusing on support tools and environments; performance and power modeling, prediction and evaluation; scheduling and load balancing; high performance architectures and compilers; data management, analytics and machine learning; cluster, cloud and edge computing; theory and algorithms for parallel and distributed processing; parallel and distributed programming,

interfaces, and languages; multicore and manycore parallelism; parallel numerical methods and applications; and accelerator computing.

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.

Software Telemetry - Jamie Riedesel 2021-08-31

Software Telemetry is a guide to operating the telemetry systems that monitor and maintain your applications. It takes a big picture view of telemetry, teaching you to manage your logging, metrics, and events as a complete end-to-end ecosystem. You'll learn the base architecture that underpins any software telemetry system, allowing you to easily integrate new systems into your existing infrastructure, and how these systems work under the hood. Throughout, you'll follow three very different companies to see

how telemetry techniques impact a greenfield startup, a large legacy enterprise, and a non-technical organization without any in-house development. You'll even cover how software telemetry is used by court processes--ensuring that when your first telemetry subpoena arrives, there's no reason to panic!

Dataspace: The Final

Frontier - Alan Sexton

2009-06-30

This book constitutes the refereed proceedings of the 26th British National Conference on Databases, BNCOD 26, held in Birmingham, UK, in July 2009. The 12 revised full papers, 2 short papers and 5 poster papers presented together with 2 keynote talks, 2 tutorial papers and summaries of 3 co-located workshops were carefully reviewed and selected from 33 submissions. The papers are organized in topical sections on data integration, warehousing and privacy; alternative data models; querying; and path queries and XML; data mining and privacy,

data integration, stream and event data processing, and query processing and optimisation.

Pro Linux System

Administration - Dennis

Matotek 2017-03-14

Implement a SOHO or SMB Linux infrastructure to expand your business and associated IT capabilities. Backed by the expertise and experienced guidance of the authors, this book provides everything you need to move your business forward. Pro Linux System Administration makes it easy for small- to medium-sized businesses to enter the world of zero-cost software running on Linux and covers all the distros you might want to use, including Red Hat, Ubuntu, Debian, and CentOS. Pro Linux System Administration takes a layered, component-based approach to open source business systems, while training system administrators as the builders of business infrastructure. Completely updated for this second edition, Dennis Matotek takes you through an infrastructure-as-

code approach, seamlessly taking you through steps along the journey of Linux administration with all you need to master complex systems. This edition now includes Jenkins, Ansible, Logstash and more. What You'll Learn: Understand Linux architecture Build, back up, and recover Linux servers Create basic networks and network services with Linux Build and implement Linux infrastructure and services including mail, web, databases, and file and print Implement Linux security Resolve Linux performance and capacity planning issues Who This Book Is For: Small to medium-sized business owners looking to run their own IT, system administrators considering migrating to Linux, and IT systems integrators looking for an extensible Linux infrastructure management approach.

Sharing Big Data Safely -

Ted Dunning 2015-09-15

Many big data-driven companies today are moving to protect certain types of data

against intrusion, leaks, or unauthorized eyes. But how do you lock down data while granting access to people who need to see it? In this practical book, authors Ted Dunning and Ellen Friedman offer two novel and practical solutions that you can implement right away. Ideal for both technical and non-technical decision makers, group leaders, developers, and data scientists, this book shows you how to: Share original data in a controlled way so that different groups within your organization only see part of the whole. You'll learn how to do this with the new open source SQL query engine Apache Drill. Provide synthetic data that emulates the behavior of sensitive data. This approach enables external advisors to work with you on projects involving data that you can't show them. If you're intrigued by the synthetic data solution, explore the log-synth program that Ted Dunning developed as open source code (available on GitHub), along with how-to instructions and tips for best practice. You'll

also get a collection of use cases. Providing lock-down security while safely sharing data is a significant challenge for a growing number of organizations. With this book, you'll discover new options to share data safely without sacrificing security.

Data Management, Analytics and Innovation - Neha Sharma
2021-10-21

This book presents the latest findings in the areas of data management and smart computing, machine learning, big data management, artificial intelligence, and data analytics, along with advances in network technologies. The book is a collection of peer-reviewed research papers presented at Fifth International Conference on Data Management, Analytics and Innovation (ICDMAI 2021), held during January 15-17, 2021, in a virtual mode. It addresses state-of-the-art topics and discusses challenges and solutions for future development. Gathering original, unpublished contributions by scientists from

around the globe, the book is mainly intended for a professional audience of researchers and practitioners in academia and industry.

Edge Computing Systems with Kubernetes - Sergio Mendez
2022-10-14

Understand how to use K3s and k3OS for different use cases and discover best practices for building an edge computing system
Key Features
A guide to implementing an edge computing environment
Reduce latency and costs for real-time applications running at the edge
Find stable and relevant cloud native open source software to complement your edge environments
Book Description
Edge computing is a way of processing information near the source of data instead of processing it on data centers in the cloud. In this way, edge computing can reduce latency when data is processed, improving the user experience on real-time data visualization for your applications. Using K3s, a light-weight Kubernetes and k3OS, a

K3s-based Linux distribution along with other open source cloud native technologies, you can build reliable edge computing systems without spending a lot of money. In this book, you will learn how to design edge computing systems with containers and edge devices using sensors, GPS modules, WiFi, LoRa communication and so on. You will also get to grips with different use cases and examples covered in this book, how to solve common use cases for edge computing such as updating your applications using GitOps, reading data from sensors and storing it on SQL and NoSQL databases. Later chapters will show you how to connect hardware to your edge clusters, predict using machine learning, and analyze images with computer vision. All the examples and use cases in this book are designed to run on devices using 64-bit ARM processors, using Raspberry Pi devices as an example. By the end of this book, you will be able to build your own edge computing

systems using the content of the chapters as Lego pieces to fit your needs. What you will learn Configure k3OS and K3s for development and production scenarios Package applications into K3s for shipped-node scenarios Deploy in occasionally connected scenarios, from one node to one million nodes Manage GitOps for applications across different locations Use open source cloud native software to complement your edge computing systems Implement observability event-driven and serverless edge applications Collect and process data from sensors at the edge and visualize it into the cloud Who this book is for This book is for engineers (developers and/or operators) seeking to bring the cloud native benefits of GitOps and Kubernetes to the edge. Anyone with basic knowledge of Linux and containers looking to learn Kubernetes using examples applied to edge computing and hardware systems will benefit from this book.

Time and Relational Theory

*Downloaded from
test.uni-cari.be.edu.doon
by guest*

- C.J. Date 2014-08-13
Time and Relational Theory provides an in-depth description of temporal database systems, which provide special facilities for storing, querying, and updating historical and future data. Traditionally, database management systems provide little or no special support for temporal data at all. This situation is changing because: Cheap storage enables retention of large volumes of historical data in data warehouses Users are now faced with temporal data problems, and need solutions Temporal features have recently been incorporated into the SQL standard, and vendors have begun to add temporal support to their DBMS products Based on the groundbreaking text Temporal Data & the Relational Model (Morgan Kaufmann, 2002) and new research led by the authors, Time and Relational Theory is the only book to offer a complete overview of the functionality of a temporal DBMS. Expert authors Nikos

Lorentzos, Hugh Darwen, and Chris Date describe an approach to temporal database management that is firmly rooted in classical relational theory and will stand the test of time. This book covers the SQL:2011 temporal extensions in depth and identifies and discusses the temporal functionality still missing from SQL. Understand how the relational model provides an ideal basis for taming the complexities of temporal databases Learn how to analyze and evaluate commercial temporal products with this timely and important information Be able to use sound principles in designing and using temporal databases Understand the temporal support recently added to SQL with coverage of the new SQL features in this unique, accurate, and authoritative reference Appreciate the benefits of a truly relational approach to the problem with this clear, user friendly presentation

Algorithms and Architectures for Parallel

Processing - Sheng Wen

2020-01-21

The two-volume set LNCS 11944-11945 constitutes the proceedings of the 19th International Conference on Algorithms and Architectures for Parallel Processing, ICA3PP 2019, held in Melbourne, Australia, in December 2019. The 73 full and 29 short papers presented were carefully reviewed and selected from 251 submissions. The papers are organized in topical sections on: Parallel and Distributed Architectures, Software Systems and Programming Models, Distributed and Parallel and Network-based Computing, Big Data and its Applications, Distributed and Parallel Algorithms, Applications of Distributed and Parallel Computing, Service Dependability and Security, IoT and CPS Computing, Performance Modelling and Evaluation.

Advanced Manufacturing and Automation XI - Yi Wang 2022

The proceedings collect selected papers from the 11th

International Workshop of Advanced Manufacturing and Automation (IWAMA 2021), held in Zhengzhou Polytechnic, China on 11 - 12 October, 2021. Topics focusing on novel techniques for manufacturing and automation in Industry 4.0 are now vital factors for the maintenance and improvement of the economy of a nation and the quality of life. It will help academic researchers and engineering to implement the concept, theory and methods in Industry 4.0 which has been a hot topic. These proceedings will make valuable contributions to academic researchers, engineers in the industry for the challenges in the 4th industry revolution and smart factories.

Traffic Monitoring and Analysis - Antonio Pescapè

2012-03-07

This book constitutes the proceedings of the 4th International Workshop on Traffic Monitoring and Analysis, TMA 2012, held in Vienna, Austria, in March 2012. The thoroughly refereed 10 full papers and 8 short

papers presented in this volume were carefully reviewed and selected from 31 submissions. The contributions are organized in topical sections on traffic analysis and characterization: new results and improved measurement techniques; measurement for QoS, security and service level agreements; and tools for network measurement and experimentation.

Streaming Architecture - Ted Dunning 2016-05-10

More and more data-driven companies are looking to adopt stream processing and streaming analytics. With this concise ebook, you'll learn best practices for designing a reliable architecture that supports this emerging big-data paradigm. Authors Ted Dunning and Ellen Friedman (Real World Hadoop) help you explore some of the best technologies to handle stream processing and analytics, with a focus on the upstream queuing or message-passing layer. To illustrate the effectiveness of these technologies, this book also

includes specific use cases.

Ideal for developers and non-technical people alike, this book describes: Key elements in good design for streaming analytics, focusing on the essential characteristics of the messaging layer New messaging technologies, including Apache Kafka and MapR Streams, with links to sample code Technology choices for streaming analytics: Apache Spark Streaming, Apache Flink, Apache Storm, and Apache Apex How stream-based architectures are helpful to support microservices Specific use cases such as fraud detection and geo-distributed data streams Ted Dunning is Chief Applications Architect at MapR Technologies, and active in the open source community. He currently serves as VP for Incubator at the Apache Foundation, as a champion and mentor for a large number of projects, and as committer and PMC member of the Apache ZooKeeper and Drill projects. Ted is on Twitter as @ted_dunning. Ellen Friedman,

a committer for the Apache Drill and Apache Mahout projects, is a solutions consultant and well-known speaker and author, currently writing mainly about big data topics. With a PhD in Biochemistry, she has years of experience as a research scientist and has written about a variety of technical topics.

Ellen is on Twitter as @Ellen_Friedman.

An Atypical ASP.NET Core 6 Design Patterns Guide - Carl-

Hugo Marcotte 2022-03-28

The professional developer's essential guide to building robust, maintainable, and flexible web apps by leveraging C# 10 and .NET 6 features and component- and application-scale design patterns

Key Features

- Apply the SOLID architectural principles and software design patterns effectively with a focus on dependency injection
- Discover modern application architectures such as vertical slice, clean architecture, and event-driven microservices
- Explore full-stack ASP.NET Core with an

overview of Blazor

Book Description

An Atypical ASP.NET Core 6 Design Patterns Guide, Second Edition approaches programming like playing with LEGO®: snapping small pieces together to create something beautiful.

Thoroughly updated for ASP.NET Core 6, with further coverage of microservices patterns, data contracts, and event-driven architecture, this book gives you the tools to build and glue reliable components together to improve your programmatic masterpieces. The chapters are organized based on scale and topic, allowing you to start small and build on a strong base, the same way that you would develop a program. You will begin by exploring basic design patterns, SOLID architectural principles, dependency injection, and other ASP.NET Core 6 mechanisms. You will explore component-scale patterns, and then move to higher level application-scale patterns and techniques to better structure your applications. Finally,

you'll advance to the client side to connect the dots with tools like Blazor and make ASP.NET Core a viable full-stack web development framework. You will supplement your learning with practical use cases and best practices, exploring a range of significant Gang of Four (GoF) design patterns along the way. By the end of the book, you will be comfortable combining and implementing patterns in different ways, and crafting software solutions of any scale. What you will learn Apply the SOLID principles for building flexible and maintainable software Get to grasp .NET dependency Injection Work with GoF design patterns such as strategy, decorator, facade, and composite Explore the MVC patterns for designing web APIs and web applications using Razor Discover layering techniques and tenets of clean architecture Become familiar with CQRS and vertical slice architecture as an alternate to layering Understand microservices and when they can benefit your

applications Build an ASP.NET user interfaces from server-side to client-side Blazor Who this book is for The book is intended for intermediate software and web developers with an understanding of .NET who want to write flexible, maintainable, and robust code for building scalable web applications. Knowledge of C# programming and an understanding of web concepts like HTTP is necessary.

Time Series Analysis with Python Cookbook - Tarek A.

Atwan 2022-06-30

Perform time series analysis and forecasting confidently with this Python code bank and reference manual Key Features Explore forecasting and anomaly detection techniques using statistical, machine learning, and deep learning algorithms Learn different techniques for evaluating, diagnosing, and optimizing your models Work with a variety of complex data with trends, multiple seasonal patterns, and irregularities Book Description Time series data is everywhere, available at

a high frequency and volume. It is complex and can contain noise, irregularities, and multiple patterns, making it crucial to be well-versed with the techniques covered in this book for data preparation, analysis, and forecasting. This book covers practical techniques for working with time series data, starting with ingesting time series data from various sources and formats, whether in private cloud storage, relational databases, non-relational databases, or specialized time series databases such as InfluxDB. Next, you'll learn strategies for handling missing data, dealing with time zones and custom business days, and detecting anomalies using intuitive statistical methods, followed by more advanced unsupervised ML models. The book will also explore forecasting using classical statistical models such as Holt-Winters, SARIMA, and VAR. The recipes will present practical techniques for handling non-stationary data, using power transforms, ACF and PACF plots, and

decomposing time series data with multiple seasonal patterns. Later, you'll work with ML and DL models using TensorFlow and PyTorch. Finally, you'll learn how to evaluate, compare, optimize models, and more using the recipes covered in the book. What you will learn Understand what makes time series data different from other data Apply various imputation and interpolation strategies for missing data Implement different models for univariate and multivariate time series Use different deep learning libraries such as TensorFlow, Keras, and PyTorch Plot interactive time series visualizations using hvPlot Explore state-space models and the unobserved components model (UCM) Detect anomalies using statistical and machine learning methods Forecast complex time series with multiple seasonal patterns Who this book is for This book is for data analysts, business analysts, data scientists, data engineers, or Python developers who want practical

Python recipes for time series analysis and forecasting techniques. Fundamental knowledge of Python programming is required. Although having a basic math and statistics background will be beneficial, it is not necessary. Prior experience working with time series data to solve business problems will also help you to better utilize and apply the different recipes in this book.

Web Database Applications with PHP and MySQL - Hugh E. Williams 2002

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

Data Mining in Time Series Databases - Mark Last
2004-06-25

Adding the time dimension to real-world databases produces Time Series Databases (TSDB) and introduces new aspects and difficulties to data mining and knowledge discovery. This book covers the state-of-the-art methodology for mining time series databases. The novel

data mining methods presented in the book include techniques for efficient segmentation, indexing, and classification of noisy and dynamic time series. A graph-based method for anomaly detection in time series is described and the book also studies the implications of a novel and potentially useful representation of time series as strings. The problem of detecting changes in data mining models that are induced from temporal databases is additionally discussed.

Contents:Segmenting Time Series: A Survey and Novel Approach (E Keogh et al.)A Survey of Recent Methods for Efficient Retrieval of Similar Time Sequences (M L Hetland)Indexing of Compressed Time Series (E Fink & K Pratt)Indexing Time-Series under Conditions of Noise (M Vlachos et al.)Change Detection in Classification Models Induced from Time Series Data (G Zeira et al.)Classification and Detection of Abnormal Events in Time Series of Graphs (H Bunke & M

Kraetzl) Boosting Interval-Based Literals: Variable Length and Early Classification (C J Alonso González & J J Rodríguez Díez) Median Strings: A Review (X Jiang et al.) Readership: Graduate students, researchers and practitioners in the fields of data mining, machine learning, databases and statistics. Keywords: Times Series; Time Series Analysis; Data Mining, Knowledge Discovery in Databases; Graphs; Graph Similarity; String Distance; Machine Learning; Segmentation; Change Detection

Computational Collective Intelligence - Ngoc Thanh Nguyen 2017-09-18

This two-volume set (LNAI 10448 and LNAI 10449) constitutes the refereed proceedings of the 9th International Conference on Collective Intelligence, ICCCI 2017, held in Nicosia, Cyprus, in September 2017. The 117 full papers presented were carefully reviewed and selected from 248 submissions. The conference focuses on the

methodology and applications of computational collective intelligence, included: multi-agent systems, knowledge engineering and semantic web, social networks and recommender systems, text processing and information retrieval, data mining methods and applications, sensor networks and internet of things, decision support & control systems, and computer vision techniques.

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.

Client-Side Data Storage - Raymond Camden 2015-12-24

One of the most useful features of today's modern browsers is the ability to store data right

on the user's computer or mobile device. Even as more people move toward the cloud, client-side storage can still save web developers a lot of time and money, if you do it right. This hands-on guide demonstrates several storage APIs in action. You'll learn how and when to use them, their pluses and minuses, and steps for implementing one or more of them in your application. Ideal for experienced web developers familiar with JavaScript, this book also introduces several open source libraries that make storage APIs easier to work with. Learn how different browsers support each client-side storage API. Work with web (aka local) storage for simple things like lists or preferences. Use IndexedDB to store nearly anything you want on the user's browser. Learn how support web apps that still use the discontinued Web SQL Database API. Explore Lockr, Dexie, and localForage, three libraries that simplify the use of storage APIs. Build a simple working application that makes

use of several storage techniques

New Trends in Databases and Information Systems -

Barbara Catania 2013-08-17

This book reports on state-of-art research and applications in the field of databases and information systems. It includes both fourteen selected short contributions, presented at the East-European Conference on Advances in Databases and Information Systems (ADBIS 2013, September 1-4, Genova, Italy), and twenty-six papers from ADBIS 2013 satellite events. The short contributions from the main conference are collected in the first part of the book, which covers a wide range of topics, like data management, similarity searches, spatio-temporal and social network data, data mining, data warehousing, and data management on novel architectures, such as graphics processing units, parallel database management systems, cloud and MapReduce environments. In contrast, the contributions from the satellite

events are organized in five different parts, according to their respective ADBIS satellite event: BiDaTA 2013 - Special Session on Big Data: New Trends and Applications); GID 2013 - The Second International Workshop on GPUs in Databases; OAIS 2013 - The Second International Workshop on Ontologies Meet Advanced Information Systems; SoBI 2013 - The First International Workshop on Social Business Intelligence: Integrating Social Content in Decision Making; and last but not least, the Doctoral Consortium, a forum for Ph.D. students. The book, which addresses academics and professionals alike, provides the readers with a comprehensive and timely overview of new trends in database and information systems research, and promotes new ideas and collaborations among the different research communities of the eastern European countries and the rest of the world.

[Creativity in Intelligent](#)

Technologies and Data Science

- Alla G. Kravets 2019-08-19

This two-volume set constitutes the proceedings of the Third Conference on Creativity in Intellectual Technologies and Data Science, CIT&DS 2019, held in Volgograd, Russia, in September 2019. The 67 full papers, 1 short paper and 3 keynote papers presented were carefully reviewed and selected from 231 submissions. The papers are organized in topical sections in the two volumes. Part I: cyber-physical systems and Big Data-driven world. Part II: artificial intelligence and deep learning technologies for creative tasks; intelligent technologies in social engineering.

Learning MySQL and MariaDB

- Russell J.T. Dyer 2015-03-30

"With an easy, step-by-step approach, this guide shows

beginners how to install, use, and maintain the world's most popular open source database: MySQL. You'll learn through real-world examples and many practical tips, including information on how to improve database performance.

Database systems such as MySQL help data handling for organizations large and small handle data, providing robust and efficient access in ways not offered by spreadsheets and other types of data stores. This book is also useful for web developers and programmers interested in adding MySQL to their skill sets. Topics include: Installation and basic administration ; Introduction to databases and SQL ; Functions, subqueries, and other query enhancements ; Improving database performance ; Accessing MySQL from popular languages"--