

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

Transaction processing is fundamental for many modern applications. These applications require the backend transaction processing engines to be available at all times as well as provide a massive horizontal scale for intensive transaction requests. Concurrency Control and Recovery features recent progress in research in online transaction processing. The book also showcases the authors' research on a highly scalable OLTP system. Its contents include the designs of an efficient multiple version storage engine, a scalable range optimistic concurrency control, high-performance Paxos-based log replication, global snapshot isolation, and fast follower recovery. This book is written for professionals, researchers, and graduate students specialising in database systems and its related fields.

Because of the explosion of unstructured data that is generated by individuals and organizations, a new storage paradigm that is called object storage has been developed. Object storage stores data in a flat namespace that scales to trillions of objects. The design of object storage also simplifies how users access data, supporting new types of applications and allowing users to access data by using various methods, including mobile devices and web applications. Data distribution and management are also simplified, allowing greater collaboration across the globe. OpenStack Swift is an emerging open source object storage software platform that is widely used for cloud storage. IBM® Spectrum Scale, which is based on IBM General Parallel File System (IBM GPFS™) technology, is a high-performance and proven product that is used to store data for thousands of mission-critical commercial installations worldwide.

Throughout this IBM Redpaper™ publication, IBM Spectrum™ Scale is used to refer to GPFS. The examples in this paper are based on IBM Spectrum Scale™ V4.2.2. IBM Spectrum Scale also automates common storage management tasks, such as tiering and archiving at scale. Together, IBM Spectrum Scale and OpenStack Swift provide an enterprise-class object storage solution that efficiently stores, distributes, and retains critical data. This paper provides instructions about setting up and configuring IBM Spectrum Scale Object Storage that is based on OpenStack Swift. It also provides an initial set of preferred practices that ensure optimal performance and reliability. This paper is intended for administrators who are familiar with IBM Spectrum Scale and OpenStack Swift components.

This book constitutes the workshop proceedings of the 16th International Conference on Database Systems for Advanced Applications, DASFAA 2011, held in Hong Kong, China, in April 2011. The volume contains six workshops, each focusing on specific research issues that contribute to the main themes of the DASFAA conference: The First International Workshop on Graph-structured Data Bases (GDB 2011); the First International Workshop on Spatial Information Modeling, Management and Mining (SIM3 2011); the International Workshop on Flash-based Database Systems (FlashDB 2011); the Second International Workshop on Social Networks and Social Media Mining on the Web (SNSMW 2011); the First International Workshop on Data Management for Emerging Network Infrastructures (DaMEN 2011); and the Fourth International

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

Workshop on Data Quality in Integration Systems (DQIS 2011).

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

What Alison's book offers over most or all of the other books on the market is that Alison is able to take a highly technical topic and present it in a manner that is easy to comprehend. It is a book that the reader will often want to read from cover to cover, but it can also act as an excellent reference. The CD includes all source code from the book, sample databases, and complete applications.

Get up to speed with core PostgreSQL tasks such as database administration, application development, database performance monitoring, and database testing Key Features Build real-world enterprise database management systems using Postgres 12 features Explore the development, administrative and security aspects of PostgreSQL 12 Implement best practices from industry experts to build powerful database applications Book Description PostgreSQL is an open-source object-relational database management system (DBMS) that provides enterprise-level services, including high performance and scalability. This book is a collection of unique projects providing you with a wealth of information relating to administering, monitoring, and testing PostgreSQL. The focus of each project is on both the development and the administrative aspects of PostgreSQL. Starting by exploring development aspects such as database design and its implementation, you'll then cover PostgreSQL administration by understanding PostgreSQL architecture, PostgreSQL performance, and high-availability clusters. Various PostgreSQL projects are explained through current technologies such as DevOps and cloud platforms using programming languages like Python and Node.js. Later, you'll get to grips with the well-known database API tool, PostgREST, before learning how to use popular PostgreSQL database testing frameworks. The book is also packed with essential tips and tricks and common patterns for working seamlessly in a production environment. All the chapters will be explained with the help of a real-world case study on a small banking application for managing ATM locations in a city. By the end of this DBMS book, you'll be proficient in building reliable database solutions as per your organization's needs. What you will learn Set up high availability PostgreSQL database clusters in the same containment, a cross-containment, and on the cloud Monitor the performance of a PostgreSQL

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

database Create automated unit tests and implement test-driven development for a PostgreSQL database Develop PostgreSQL apps on cloud platforms using DevOps with Python and Node.js Write robust APIs for PostgreSQL databases using Python programming, Node.js, and PostgREST Create a geospatial database using PostGIS and PostgreSQL Implement automatic configuration by Ansible and Terraform for Postgres Who this book is for This PostgreSQL book is for database developers, database administrators, data architects, or anyone who wants to build end-to-end database projects using Postgres. This book will also appeal to software engineers, IT technicians, computer science researchers, and university students who are interested in database development and administration. Some familiarity with PostgreSQL and Linux is required to grasp the concepts covered in the book effectively.

Prepare for Microsoft Exam 70-486—and help demonstrate your real-world mastery of developing ASP.NET MVC-based solutions. Designed for experienced developers ready to advance their status, Exam Ref focuses on the critical-thinking and decision-making acumen needed for success at the Microsoft Specialist level. Focus on the expertise measured by these objectives: Design the application architecture Design the user experience Develop the user experience Troubleshoot and debug web applications Design and implement security This Microsoft Exam Ref: Organizes its coverage by exam objectives. Features strategic, what-if scenarios to challenge you. The purpose of this book is to provide an up-to-date and systematic introduction to the principles and algorithms of machine learning. The definition of learning is broad enough to include most tasks that we commonly call “learning” tasks, as we use the word in daily life. It is also broad enough to encompass computers that improve from experience in quite straightforward ways. The book will be of interest to industrial engineers and scientists as well as academics who wish to pursue machine learning. The book is intended for both graduate and postgraduate students in fields such as computer science, cybernetics, system sciences, engineering, statistics, and social sciences, and as a reference for software professionals and practitioners. The wide scope of the book provides a good introduction to many approaches of machine learning, and it is also the source of useful bibliographical information.

Visualization and analysis tools, techniques, and algorithms have undergone a rapid evolution in recent decades to accommodate explosive growth in data size and complexity and to exploit emerging multi- and many-core computational platforms. High Performance Visualization: Enabling Extreme-Scale Scientific Insight focuses on the subset of scientific visualization concerned with algorithm design, implementation, and optimization for use on today’s largest computational platforms. The book collects some of the most seminal work in the field, including algorithms and implementations running at the highest levels of concurrency and used by scientific researchers worldwide. After introducing the fundamental concepts of parallel visualization, the book explores approaches to accelerate visualization and analysis operations on high performance computing platforms. Looking to the future and anticipating changes to computational platforms in the transition from the petascale to exascale regime, it presents the main research challenges and describes several contemporary, high performance visualization implementations. Reflecting major concepts in high performance visualization, this book unifies a large and diverse body of computer science research, development, and practical applications. It describes the state of the art at the intersection of scientific visualization, large data, and high performance computing trends, giving readers the foundation to apply the concepts and carry out future research in this area.

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

This is the only book on the market to focus on addressing issues of building highly scalable database applications with .NET technologies. Comprehensive coverage includes building .NET applications for all the major RDBMSs: SQL Server, Oracle, DB2, and MySQL. This book examines the field of parallel database management systems and illustrates the great variety of solutions based on a shared-storage or a shared-nothing architecture. Constantly dropping memory prices and the desire to operate with low-latency responses on large sets of data paved the way for main memory-based parallel database management systems. However, this area is currently dominated by the shared-nothing approach in order to preserve the in-memory performance advantage by processing data locally on each server. The main argument this book makes is that such an unilateral development will cease due to the combination of the following three trends: a) Today's network technology features remote direct memory access (RDMA) and narrows the performance gap between accessing main memory on a server and of a remote server to and even below a single order of magnitude. b) Modern storage systems scale gracefully, are elastic and provide high-availability. c) A modern storage system such as Stanford's RAM Cloud even keeps all data resident in the main memory. Exploiting these characteristics in the context of a main memory-based parallel database management system is desirable. The book demonstrates that the advent of RDMA-enabled network technology makes the creation of a parallel main memory DBMS based on a shared-storage approach feasible.

This book constitutes the refereed proceedings of the 15th International Conference on Parallel Computing, Euro-Par 2009, held in Delft, The Netherlands, in August 2009. The 85 revised papers presented were carefully reviewed and selected from 256 submissions. The papers are organized in topical sections on support tools and environments; performance prediction and evaluation; scheduling and load balancing; high performance architectures and compilers; parallel and distributed databases; grid, cluster, and cloud computing; peer-to-peer computing; distributed systems and algorithms; parallel and distributed programming; parallel numerical algorithms; multicore and manycore programming; theory and algorithms for parallel computation; high performance networks; and mobile and ubiquitous computing.

In the last two decades, the biannual ECPPM (European Conference on Product and Process Modelling) conference series has provided a unique platform for the presentation and discussion of the most recent advances with regard to the ICT (Information and Communication Technology) applications in the AEC/FM (Architecture, Engineering, Construction and

As open systems continue to replace traditional mainframe systems, system scalability is becoming an increasingly important topic. This guide offers techniques for designing reliable and scalable online transaction processing (OLTP) applications using Oracle. It covers hardware and I/O operation; benchmark and database monitoring systems; Oracle internals, operation, and implementation; and UNIX operating system issues that impact Oracle performance and scalability. The CD-ROM contains source code for dbaman, code examples, and public domain software. Annotation copyrighted by Book News, Inc., Portland, OR

As the age of Big Data emerges, it becomes necessary to take the five dimensions of Big Data- volume, variety, velocity, volatility, and veracity- and focus these dimensions towards one critical emphasis - value. The Encyclopedia of Business Analytics and Optimization confronts the challenges of information retrieval in the age of Big Data by exploring recent advances in the areas of knowledge management, data visualization, interdisciplinary communication, and others. Through its critical approach and practical application, this book will be a must-have reference for any professional, leader, analyst, or manager interested in making the most of the knowledge resources at their disposal.

The aims of these proceedings are to provide a complete coverage of the areas outlined, and to bring together researchers from academic and industry to share ideas, challenges, and

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

solutions relating to the multifaceted aspects of this field. New multimedia standards (for example, MPEG-21) facilitate the seamless integration of multiple modalities into interoperable multimedia frameworks, transforming the way people work and interact with multimedia data. These key technologies and multimedia solutions interact and collaborate with each other in increasingly effective ways, contributing to the multimedia revolution and having a significant impact across a wide spectrum of consumer, business, healthcare, education, and governmental domains.

First complete book in the market to quickly get developers up-to-speed with VB.NET.

Give users the real-time experience they expect, by using Elixir and Phoenix Channels to build applications that instantly react to changes and reflect the application's true state. Learn how Elixir and Phoenix make it easy and enjoyable to create real-time applications that scale to a large number of users. Apply system design and development best practices to create applications that are easy to maintain. Gain confidence by learning how to break your applications before your users do. Deploy applications with minimized resource use and maximized performance. Real-time applications come with real challenges - persistent connections, multi-server deployment, and strict performance requirements are just a few. Don't try to solve these challenges by yourself - use a framework that handles them for you. Elixir and Phoenix Channels provide a solid foundation on which to build stable and scalable real-time applications. Build applications that thrive for years to come with the best-practices found in this book. Understand the magic of real-time communication by inspecting the WebSocket protocol in action. Avoid performance pitfalls early in the development lifecycle with a catalog of common problems and their solutions. Leverage GenStage to build a data pipeline that improves scalability. Break your application before your users do and confidently deploy them. Build a real-world project using solid application design and testing practices that help make future changes a breeze. Create distributed apps that can scale to many users with tools like Phoenix Tracker. Deploy and monitor your application with confidence and reduce outages. Deliver an exceptional real-time experience to your users, with easy maintenance, reduced operational costs, and maximized performance, using Elixir and Phoenix Channels. What You Need: You'll need Elixir 1.9+ and Erlang/OTP 22+ installed on a Mac OS X, Linux, or Windows machine.

This book highlights the different types of data architecture and illustrates the many possibilities hidden behind the term "Big Data", from the usage of No-SQL databases to the deployment of stream analytics architecture, machine learning, and governance. Scalable Big Data Architecture covers real-world, concrete industry use cases that leverage complex distributed applications, which involve web applications, RESTful API, and high throughput of large amount of data stored in highly scalable No-SQL data stores such as Couchbase and Elasticsearch. This book demonstrates how data processing can be done at scale from the usage of NoSQL datastores to the combination of Big Data distribution. When the data processing is too complex and involves different

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

processing topology like long running jobs, stream processing, multiple data sources correlation, and machine learning, it's often necessary to delegate the load to Hadoop or Spark and use the No-SQL to serve processed data in real time. This book shows you how to choose a relevant combination of big data technologies available within the Hadoop ecosystem. It focuses on processing long jobs, architecture, stream data patterns, log analysis, and real time analytics. Every pattern is illustrated with practical examples, which use the different open source projects such as Logstash, Spark, Kafka, and so on. Traditional data infrastructures are built for digesting and rendering data synthesis and analytics from large amount of data. This book helps you to understand why you should consider using machine learning algorithms early on in the project, before being overwhelmed by constraints imposed by dealing with the high throughput of Big data. Scalable Big Data Architecture is for developers, data architects, and data scientists looking for a better understanding of how to choose the most relevant pattern for a Big Data project and which tools to integrate into that pattern.

Storage Systems: Organization, Performance, Coding, Reliability and Their Data Processing covers the coding, reliability and performance of popular RAID organizations: RAID1 mirrored disks, RAID5/6/7 1/2/3-disk failure tolerant - 1/2/3DFT arrays. Readers will learn about the storage of files, SQL and NoSQL databases on disk and SSD to achieve higher efficiency. As data compression, deduplication and encryption techniques for storage systems have led to new technologies, startups and techniques to save power in storage and server systems, the book discusses Fast Array of Wimpy Nodes (FAWN) at CMU, RAMCloud at Stanford, and key-value flash Lightstore at MIT, along with several storage proposals. Finally, storage technologies from punched cards up to ash memories and beyond are discussed, along with the data placement and the scheduling of magnetic disks. Provides readers with an in-depth understanding of the architecture and operation of computer storage systems Includes descriptions of various RAID levels, their coding, organization, performance and reliability Covers techniques for efficient and secure data storage through data compression, deduplication and encryption Presents readers with an in-depth understanding of the storage of files and SQL and NoSQL databases

This IBM® Redpaper™ publication gives readers a broad understanding of IBM Bluemix™ cloud application development platform capabilities. Providing a platform as a service (PaaS) environment as one of its run times, along with containers and virtual machines, Bluemix uses the Cloud Foundry project as one of its open source technologies to accelerate new application development and DevOps methods. It provides optimized and flexible workloads, enables continuous availability, and simplifies delivery and manageability of an application by providing prebuilt services and hosting capabilities. The paper reviews the Bluemix architecture, explains how it works, describes key concepts and components, and provides an overview of Bluemix security. It also covers the

various Bluemix service categories and the services within each category. This information will help anyone who is interested in exploring the potential and capabilities of Bluemix and its services.

This book constitutes the workshop proceedings of the 23rd International Conference on Database Systems for Advanced Applications, DASFAA 2018, held in Gold Coast, QLD, Australia, in May 2018. The 23 full papers presented were carefully selected and reviewed from 44 submissions to the four following workshops: the 5th International Workshop on Big Data Management and Service, BDMS 2018; the Third International Workshop on Big Data Quality Management, BDQM 2018; the Second International Workshop on Graph Data Management and Analysis, GDMA 2018; and the 5th International Workshop on Semantic Computing and Personalization, SeCoP 2018.

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

The two-volume set LNCS 9366 and 9367 constitutes the refereed proceedings of the 14th International Semantic Web Conference, ISWC 2015, held in Bethlehem, PA, USA, in October 2015. The International Semantic Web Conference is the premier forum for Semantic Web research, where cutting edge scientific results and technological innovations are presented, where problems and solutions are discussed, and where the future of this vision is being developed. It brings together specialists in fields such as artificial intelligence, databases, social networks, distributed computing, Web engineering, information systems, human-computer interaction, natural language processing, and the

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

social sciences. The papers cover topics such as querying with SPARQL; querying linked data; linked data; ontology-based data access; ontology alignment; reasoning; instance matching, entity resolution and topic generation; RDF data dynamics; ontology extraction and generation; knowledge graphs and scientific data publication; ontology instance alignment; knowledge graphs; data processing, IoT, sensors; archiving and publishing scientific data; IoT and sensors; experiments; evaluation; and empirical studies. Part 1 (LNCS 9366) contains a total of 38 papers which were presented in the research track. They were carefully reviewed and selected from 172 submissions. Part 2 (LNCS 9367) contains 14 papers from the in-use and software track, 8 papers from the datasets and ontologies track, and 7 papers from the empirical studies and experiments track, selected, respectively, from 33, 35, and 23 submissions. 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.

All applications use data, and most applications also need to store this data somewhere. In the world of business solutions, this often meant creating a relational database. However, relational technology is not always the best solution to meet the increasingly complex data-processing requirements of modern business systems, especially when this processing involves storing and retrieving massive amounts of data. The advent of NoSQL databases has changed the way in which organizations have started to think about the way in which they structure their data. There is no standard definition of what a NoSQL database is other than they are all non-relational. They are less generalized than relational databases, but the driving force behind most NoSQL databases is focused efficiency and high scalability. The downside of NoSQL is that no single database is likely to be able to support the complete range of business requirements mandated by your applications. How do you select the most appropriate database to use, or should you remain with the relational model? A modern business application is not restricted to using a single data store, and an increasing number of solutions are now based on a polyglot architecture. The key to designing a successful application is to understand which databases best meet the needs of the various parts of the system, and how to combine these databases into a single, seamless solution. This guide helps you understand these challenges and enables you to apply the principles of NoSQL databases and polyglot solutions in your own environment. To help illustrate how to build a polyglot solution, this guide presents a case study of a fictitious company faced with building a highly scalable web application capable of supporting many thousands of concurrent users.

* Explains how to plan and implement scalable application designs using .NET 2.0 and both traditional and Service Oriented (SOA) models. * Gives detailed overview advice across the whole breadth of the project from platform to application to database in order to provide a comprehensive treatment rather than just focus on one particular issue. * Uses the very latest version of Web Services Enhancements (WSE 3.0) when most

competing titles still use the previous version (WSE 2.0 or WSE 1.0) and includes detailed consideration of the new Windows Server System and advises how to select the correct setup for your project.⁵

Embedded and ubiquitous computing systems have considerably increased their scope of application over the past few years, and they now also include mission- and business-critical scenarios. The advances call for a variety of compelling issues, including dependability, real-time, quality-of-service, autonomy, resource constraints, seamless interaction, middleware support, modeling, verification, validation, etc. The International Workshop on Software Technologies for Future Embedded and Ubiquitous Systems (SEUS) brings together experts in the field of embedded and ubiquitous computing systems with the aim of exchanging ideas and advancing the state of the art about the above-mentioned issues. I was honored to chair the sixth edition of the workshop, which continued the tradition of past editions with high-quality research results. I was particularly pleased to host the workshop in the wonderful scenario of Capri, with its stunning views and traditions. The workshop started in 2003 as an IEEE event, and then in 2007 it became a flagship event of the IFIP Working Group 10.2 on embedded systems. The last few editions, held in Hakodate (Japan), Vienna (Austria), Seattle (USA), Gyeongju (Korea), and Santorini (Greece), were co-located with the International Symposium on Object/Component/Service-Oriented Real-Time Distributed Computing (ISORC). This year, SEUS was held as a stand-alone event for the first time, and, despite the additional organizational difficulties, it resulted in a high-quality event, with papers from four continents (from USA, Europe, East Asia and Australia), (co-) authored and presented from senior scientists coming from academia or leading industrial research centers.

Since 2008, Globe has been an annual international conference on data management in grid and peer-to-peer systems. Initially, grid and peer-to-peer systems experienced significant success in scientific and file sharing applications. Today, these systems cover the management of large, distributed and heterogeneous data. These systems are characterized by high heterogeneity, high autonomy and dynamics of nodes, decentralization of control and large-scale distribution of resources. Research on data management in grid and peer-to-peer, relatively recent, aims to scale distributed systems and applications that require effective management of voluminous, large-scale distributed and heterogeneous data. The third edition of the international conference Globe was held in Bilbao, Spain during September 1-2, 2010. Globe provided opportunities for academia or industry researchers to present and discuss the latest research and applications on data management in grid and peer-to-peer systems. Globe 2010 received 26 papers from 15 countries. The reviewing process led to the acceptance of 13 papers for presentation at the conference and inclusion in this LNCS volume. Each paper was reviewed by at least two Program Committee members. The conference would not have been possible without the support of the Program Committee members, external reviewers, Organizing Committee, members of the DEXA conference and the authors. In particular, we would like to thank Gabriela Wagner and Roland Wagner (FAW, University of Linz) for their help in the realization of this conference.

This book constitutes the thoroughly refereed post-conference proceedings of the International Conference on Scalable Information Systems, INFOSCALE 2014, held in September 2014 in Seoul, South Korea. The 9 revised full papers presented were

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

carefully reviewed and selected from 14 submissions. The papers cover a wide range of topics such as scalable data analysis and big data applications.

Data Access for Highly-Scalable Solutions Using SQL, NoSQL, and Polyglot Persistence Microsoft patterns & practices

- * Describes the architecture of a scalable .NET application using various Microsoft technologies not only .NET but also SQL Server 2000.
- * Focuses the importance of correct design to avoid scalability problems in production.
- * Gives a thorough overview of scalability design suitable for IT Architects, system designers and developers.
- * Teaches the essential application frameworks to enhance scalability in a multi tiered application.

The two volume set LNCS 6443 and LNCS 6444 constitutes the proceedings of the 17th International Conference on Neural Information Processing, ICONIP 2010, held in Sydney, Australia, in November 2010. The 146 regular session papers presented were carefully reviewed and selected from 470 submissions. The papers of part I are organized in topical sections on neurodynamics, computational neuroscience and cognitive science, data and text processing, adaptive algorithms, bio-inspired algorithms, and hierarchical methods. The second volume is structured in topical sections on brain computer interface, kernel methods, computational advance in bioinformatics, self-organizing maps and their applications, machine learning applications to image analysis, and applications.

Scaling Java enterprise applications beyond just programming techniques--this is the next level. This volume covers all the technologies Java developers need to build scalable, high-performance Web applications. The book also covers servlet-based session management, EJB application logic, database design and integration, and more.

Born after World War II, large-scale experimental high-energy physics (HEP) has found itself limited ever since by available accelerator, detector and computing technologies. Accordingly, HEP has made significant contributions to the development of these fields, more often than not driving their innovations. The invention of the World Wide Web at CERN is merely the best-known example out of many. This book is the first comprehensive account to trace the history of this pioneering spirit in the field of computing technologies. It covers everything up to and including the present-day handling of the huge demands imposed upon grid and distributed computing by full-scale LHC operations—operations which have for years involved many thousands of collaborating members worldwide and accordingly provide the original and natural testbed for grid computing concepts. This book takes the reader on a guided tour encompassing all relevant topics, including programming languages, software engineering, large databases, the Web, and grid- and cloud computing. The important issue of intellectual property regulations for distributed software engineering and computing is also addressed. Aply, the book closes with a visionary chapter of what may lie ahead. Approachable and requiring only basic understanding of physics and computer sciences, this book is intended for both education and research.

Grid Computing and Cluster Computing are advanced topics and latest trends in computer science that find a place in the computer science and information technology curricula of many engineering institutes and universities today. Divided into two parts—Part I, Grid Computing and Part II, Cluster Computing—, this compact and concise text strives to make the concepts of grid computing and cluster computing comprehensible to the students through its fine presentation and accessible style. Part I of the book enables the student not only to understand the concepts involved in grid computing but also to build their own grids for specific applications. Similarly, as today supercomputers are being built using cluster computing architectures, Part II provides an insight into the basic principles involved in cluster computing and equips the readers with the knowledge to build their own clusters in-house. Diagrams are used to illustrate

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

the concepts discussed and to enable the reader to actually construct a grid or a cluster himself. The book is intended as a text for undergraduate and postgraduate students of computer science and engineering, information technology (B.Tech./M.Tech. Computer Science and Engineering/IT), and post-graduate students of computer science/information technology (M.Sc. Computer Science and M.Sc. IT). Besides, practising engineers and computer science professionals should find the text very useful.

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

[Copyright: 553ee20da8d32853e88ce847acf1669f](#)