

Get Free Accelerating Sql Database Operations On A Gpu With Cuda

Accelerating Sql Database Operations On A Gpu With Cuda | 6957a634b43137280fa88731baf1b19b

Advances in Databases and Information Systems Database and Expert Systems Applications Storage Systems Architecting and Deploying DB2 with BLU Acceleration Complete Analytics with IBM DB2 Query Management Facility: Accelerating Well-Informed Decisions Across the Enterprise Motion in Games GPU Computing and Applications Data Management Technologies and Applications Multi-Processor System-on-Chip 2 Information Search, Integration, and Personalization Microsoft Azure Infrastructure Services for Architects Advances in GPU Research and Practice New Trends in Databases and Information Systems Enabling Real-time Analytics on IBM z Systems Platform Big Data Analytics and Knowledge Discovery Beyond Databases, Architectures and Structures. Facing the Challenges of Data Proliferation and Growing Variety Advances in Computer Communication and Computational Sciences Accelerating MATLAB Performance High Performance Computing Information and Communication Technology for Intelligent Systems Transformation von Multiphysics-Modellen in einen FPGA-Entwurf für den echtzeitfähigen HiL-Test eingebetteter Systeme Proceedings of the Second International Conference on Computational Intelligence and Informatics Proceedings of ELM-2017 Web-Age Information Management Information and Communications Security Reliability and Performance with IBM DB2 Analytics Accelerator V4.1 Big Data Management and Processing Databases Theory and Applications Query Acceleration for Business Using IBM Informix Warehouse Accelerator InfoWorld Accelerating Digital Transformation on Z Using Data Virtualization Computer, Intelligent Computing and Education Technology Expert Oracle RAC 12c Wescon/89 Conference Record Beyond Databases, Architectures and Structures. Advanced Technologies for Data Mining and Knowledge Discovery 2012 International Conference on Information Technology and Management Science (ICITMS 2012) Proceedings Advanced Informatics for Computing Research Accelerating Data Transformation with IBM DB2 Analytics Accelerator for z/OS Transactions on Large-Scale Data- and Knowledge-Centered Systems XVEncyclopedia of Business Analytics and Optimization

This book constitutes the refereed proceedings of the 15th International Conference on Advances in Databases and Information Systems, ADBIS 2011, held in Vienna, Austria, in September 2011. The 30 revised full papers presented together with 2 full length invited talks were carefully reviewed and selected from 105 submissions. They are organized in topical sections on query processing; data warehousing; DB systems; spatial data; information systems; physical DB design; evolution, integrity, security; and data semantics.

Regarding online transaction processing (OLTP) workloads, IBM® z Systems™ platform, with IBM DB2®, data sharing, Workload Manager (WLM), geoplex, and other high-end features, is the widely acknowledged leader. Most customers now integrate business analytics with OLTP by running, for example, scoring functions from transactional context for real-time analytics or by applying machine-learning algorithms on enterprise data that is kept on the mainframe. As a result, IBM adds investment so clients can keep the complete lifecycle for data analysis, modeling, and scoring on z Systems control in a cost-efficient way, keeping the qualities of services in availability, security, reliability that z Systems

Get Free Accelerating Sql Database Operations On A Gpu With Cuda

solutions offer. Because of the changed architecture and tighter integration, IBM has shown, in a customer proof-of-concept, that a particular client was able to achieve an orders-of-magnitude improvement in performance, allowing that client's data scientist to investigate the data in a more interactive process. Open technologies, such as Predictive Model Markup Language (PMML) can help customers update single components instead of being forced to replace everything at once. As a result, you have the possibility to combine your preferred tool for model generation (such as SAS Enterprise Miner or IBM SPSS® Modeler) with a different technology for model scoring (such as Zementis, a company focused on PMML scoring). IBM SPSS Modeler is a leading data mining workbench that can apply various algorithms in data preparation, cleansing, statistics, visualization, machine learning, and predictive analytics. It has over 20 years of experience and continued development, and is integrated with z Systems. With IBM DB2 Analytics Accelerator 5.1 and SPSS Modeler 17.1, the possibility exists to do the complete predictive model creation including data transformation within DB2 Analytics Accelerator. So, instead of moving the data to a distributed environment, algorithms can be pushed to the data, using cost-efficient DB2 Accelerator for the required resource-intensive operations. This IBM Redbooks® publication explains the overall z Systems architecture, how the components can be installed and customized, how the new IBM DB2 Analytics Accelerator loader can help efficient data loading for z Systems data and external data, how in-database transformation, in-database modeling, and in-transactional real-time scoring can be used, and what other related technologies are available. This book is intended for technical specialists and architects, and data scientists who want to use the technology on the z Systems platform. Most of the technologies described in this book require IBM DB2 for z/OS®. For acceleration of the data investigation, data transformation, and data modeling process, DB2 Analytics Accelerator is required. Most value can be achieved if most of the data already resides on z Systems platforms, although adding external data (like from social sources) poses no problem at all.

This IBM® Redpaper™ publication introduces a new data virtualization capability that enables IBM z/OS® data to be combined with other enterprise data sources in real-time, which allows applications to access any live enterprise data anytime and use the power and efficiencies of the IBM Z® platform. Modern businesses need actionable and timely insight from current data. They cannot afford the time that is necessary to copy and transform data. They also cannot afford to secure and protect each copy of personally identifiable information and corporate intellectual property. Data virtualization enables direct connections to be established between multiple data sources and the applications that process the data. Transformations can be applied, in line, to enable real-time access to data, which opens up many new ways to gain business insight with less IT infrastructure necessary to achieve those goals. Data virtualization can become the backbone for advanced analytics and modern applications. The IBM Data Virtualization Manager for z/OS (DVM) can be used as a stand-alone product or as a utility that is used by other products. Its goal is to enable access to live mainframe transaction data and make it usable by any application. This “this what?” enables customers to use the strengths of mainframe processing with new agile applications. Additionally, its modern development environment and code-generating capabilities enable any developer to update, access, and combine mainframe data easily by using modern APIs and languages. If data is the foundation for building new insights, IBM DVM is a key tool for providing easy, cost-efficient access to that foundation.

Get Free Accelerating Sql Database Operations On A Gpu With Cuda

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.

Advances in GPU Research and Practice focuses on research and practices in GPU based systems. The topics treated cover a range of issues, ranging from hardware and architectural issues, to high level issues, such as application systems, parallel programming, middleware, and power and energy issues. Divided into six parts, this edited volume provides the latest research on GPU computing. Part I: Architectural Solutions focuses on the architectural topics that improve on performance of GPUs, Part II: System Software discusses OS, compilers, libraries, programming environment, languages, and paradigms that are proposed and analyzed to help and support GPU programmers. Part III: Power and Reliability Issues covers different aspects of energy, power, and reliability concerns in GPUs. Part IV: Performance Analysis illustrates mathematical and analytical techniques to predict different performance metrics in GPUs. Part V: Algorithms presents how to design efficient algorithms and analyze their complexity for GPUs. Part VI: Applications and Related Topics provides use cases and examples of how GPUs are used across many sectors. Discusses how to maximize power and obtain peak reliability when designing, building, and using GPUs Covers system software (OS, compilers), programming environments, languages, and paradigms proposed to help and support GPU programmers Explains how to use mathematical and analytical techniques to predict different performance metrics in GPUs Illustrates the design of efficient GPU algorithms in areas such as bioinformatics, complex systems, social networks, and cryptography Provides applications and use case scenarios in several different verticals, including medicine, social sciences, image processing, and telecommunications

From the Foreword: "Big Data Management and Processing is [a] state-of-the-art book that deals with a wide range of topical themes in the field of Big Data. The book, which probes many issues related to this exciting and rapidly growing field, covers processing, management, analytics, and applications [It] is a very valuable addition to the literature. It will serve as a source of up-to-date research in this continuously developing area. The book also provides an opportunity for researchers to explore the use of advanced computing technologies and their impact on enhancing our capabilities to conduct more sophisticated studies." ---Sartaj Sahni, University of Florida, USA "Big Data Management and Processing covers the latest Big Data research results in processing, analytics, management and applications. Both fundamental insights and representative applications are provided. This book is a timely and valuable resource for students, researchers and seasoned practitioners in Big Data fields. --Hai Jin, Huazhong University of Science and Technology, China Big Data Management and Processing explores a range of big data related issues and their impact on the design of new computing systems. The twenty-one chapters were carefully selected and feature contributions from several outstanding researchers. The book endeavors to strike a balance between theoretical and practical coverage of innovative problem solving techniques for a range of platforms. It serves as a repository of paradigms, technologies, and applications that target different facets of big data computing systems. The first part of the book explores energy and resource management issues, as well as legal compliance and quality management for Big Data. It covers In-Memory

Get Free Accelerating Sql Database Operations On A Gpu With Cuda

computing and In-Memory data grids, as well as co-scheduling for high performance computing applications. The second part of the book includes comprehensive coverage of Hadoop and Spark, along with security, privacy, and trust challenges and solutions. The latter part of the book covers mining and clustering in Big Data, and includes applications in genomics, hospital big data processing, and vehicular cloud computing. The book also analyzes funding for Big Data projects.

This book constitutes the refereed proceedings of the 29th Australasian Database Conference, ADC 2018, held in Gold Coast, QLD, Australia, in May 2018. The 23 full papers plus 6 short papers presented together with 3 demo papers were carefully reviewed and selected from 53 submissions. The Australasian Database Conference is an annual international forum for sharing the latest research advancements and novel applications of database systems, data-driven applications, and data analytics between researchers and practitioners from around the globe, particularly Australia and New Zealand.

This book presents a collection of state of the art research on GPU Computing and Application. The major part of this book is selected from the work presented at the 2013 Symposium on GPU Computing and Applications held in Nanyang Technological University, Singapore (Oct 9, 2013). Three major domains of GPU application are covered in the book including (1) Engineering design and simulation; (2) Biomedical Sciences; and (3) Interactive & Digital Media. The book also addresses the fundamental issues in GPU computing with a focus on big data processing. Researchers and developers in GPU Computing and Applications will benefit from this book. Training professionals and educators can also benefit from this book to learn the possible application of GPU technology in various areas.

This book constitutes the refereed proceedings of the 14th International Conference entitled Beyond Databases, Architectures and Structures, BDAS 2018, held in Poznań, Poland, in September 2018, during the IFIP World Computer Congress. It consists of 38 carefully reviewed papers selected from 102 submissions. The papers are organized in topical sections, namely big data and cloud computing; architectures, structures and algorithms for efficient data processing; artificial intelligence, data mining and knowledge discovery; text mining, natural language processing, ontologies and semantic web; image analysis and multimedia mining.

The volume contains 69 high quality papers presented at International Conference on Computational Intelligence and Informatics (ICCI 2017). The conference was held during 25-27, September, 2017 at Department of Computer Science and Engineering, JNTUHCHE, Hyderabad, Telangana, India. This volume contains papers mainly focused on data mining, wireless sensor networks, parallel computing, image processing, network security, MANETS, natural language processing, and internet of things.

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

This two-volume set, LNCS 12923 and 12924, constitutes the thoroughly refereed proceedings of the 5th International Conference on Database and Expert Systems Applications, DEXA 2021. Due to COVID-19 pandemic, the conference was held

Get Free Accelerating Sql Database Operations On A Gpu With Cuda

virtually. The 37 full papers presented together with 31 short papers in these volumes were carefully reviewed and selected from a total of 149 submissions. The papers are organized around the following topics: big data; data analysis and data modeling; data mining; databases and data management; information retrieval; prediction and decision support.

The IBM® DB2® Analytics Accelerator for IBM z/OS® is a high-performance appliance that integrates the IBM zEnterprise® infrastructure with IBM PureData™ for Analytics, powered by IBM Netezza® technology. With this integration, you can accelerate data-intensive and complex queries in a DB2 for z/OS highly secure and available environment. DB2 and the Analytics Accelerator appliance form a self-managing hybrid environment running online transaction processing and online transactional analytical processing concurrently and efficiently. These online transactions run together with business intelligence and online analytic processing workloads. DB2 Analytics Accelerator V4.1 expands the value of high-performance analytics. DB2 Analytics Accelerator V4.1 opens to static Structured Query Language (SQL) applications and row set processing, minimizes data movement, reduces latency, and improves availability. This IBM Redbooks® publication provides technical decision-makers with an understanding of the benefits of version 4.1 of the Analytics Accelerator with DB2 11 for z/OS. It describes the installation of the new functions, and the advantages to existing analytical processes as measured in our test environment. This book also introduces the DB2 Analytics Accelerator Loader V1.1, a tool that facilitates the data population of the DB2 Analytics Accelerator.

IBM® Informix® Warehouse Accelerator is a state-of-the-art in-memory database that uses affordable innovations in memory and processor technology and trends in novel ways to boost query performance. It is a disruptive technology that changes how organizations provide analytics to its operational and historical data. Informix Warehouse Accelerator uses columnar, in-memory approach to accelerate even the most complex warehouse and operational queries without application changes or tuning. This IBM Redbooks® publication provides a comprehensive look at the technology and architecture behind the system. It contains information about the tools, data synchronization, and query processing capabilities of Informix Warehouse Accelerator, and provides steps to implement data analysis by using Informix Warehouse Accelerator within an organization. This book is intended for IBM Business Partners and clients who are looking for low-cost solutions to boost data warehouse query performance.

Annotation. This book constitutes the refereed proceedings of the 12th International Conference on Information and Communications Security, ICICS 2010, held in Barcelona, Spain, in December 2010. The 31 revised full papers presented together with an invited talk were carefully reviewed and selected from 135 submissions. The papers are organized in topical sections on access control, public key cryptography and cryptanalysis, security in distributed and mobile systems, cryptanalysis, authentication, fair exchange protocols, anonymity and privacy, software security, proxy cryptosystems, and intrusion detection systems.

IBM® DB2® with BLU Acceleration is a revolutionary technology that is delivered in DB2 for Linux, UNIX, and Windows Release 10.5. BLU Acceleration delivers breakthrough performance improvements for analytic queries by using dynamic in-memory columnar technologies. Different from other vendor solutions, BLU Acceleration allows the unified computing of OLTP and analytics data inside a single database, therefore, removing barriers and accelerating results for users. With

Get Free Accelerating Sql Database Operations On A Gpu With Cuda

observed hundredfold improvement in query response time, BLU Acceleration provides a simple, fast, and easy-to-use solution for the needs of today's organizations; quick access to business answers can be used to gain a competitive edge, lower costs, and more. This IBM Redbooks® publication introduces the concepts of DB2 with BLU Acceleration. It discusses the steps to move from a relational database to using BLU Acceleration, optimizing BLU usage, and deploying BLU into existing analytic solutions today, with an example of IBM Cognos®. This book also describes integration of DB2 with BLU Acceleration into SAP Business Warehouse (SAP BW) and SAP's near-line storage solution on DB2. This publication is intended to be helpful to a wide-ranging audience, including those readers who want to understand the technologies and those who have planning, deployment, and support responsibilities.

The MATLAB® programming environment is often perceived as a platform suitable for prototyping and modeling but not for "serious" applications. One of the main complaints is that MATLAB is just too slow. Accelerating MATLAB Performance aims to correct this perception by describing multiple ways to greatly improve MATLAB program speed. Packed with thousands of helpful tips, it leaves no stone unturned, discussing every aspect of MATLAB. Ideal for novices and professionals alike, the book describes MATLAB performance in a scale and depth never before published. It takes a comprehensive approach to MATLAB performance, illustrating numerous ways to attain the desired speedup. The book covers MATLAB, CPU, and memory profiling and discusses various tradeoffs in performance tuning. It describes both the application of standard industry techniques in MATLAB, as well as methods that are specific to MATLAB such as using different data types or built-in functions. The book covers MATLAB vectorization, parallelization (implicit and explicit), optimization, memory management, chunking, and caching. It explains MATLAB's memory model and details how it can be leveraged. It describes the use of GPU, MEX, FPGA, and other forms of compiled code, as well as techniques for speeding up deployed applications. It details specific tips for MATLAB GUI, graphics, and I/O. It also reviews a wide variety of utilities, libraries, and toolboxes that can help to improve performance. Sufficient information is provided to allow readers to immediately apply the suggestions to their own MATLAB programs. Extensive references are also included to allow those who wish to expand the treatment of a particular topic to do so easily. Supported by an active website, and numerous code examples, the book will help readers rapidly attain significant reductions in development costs and program run times.

This book constitutes the thoroughly refereed proceedings of the Fourth International Conference on Data Technologies and Applications, DATA 2016, held in Colmar, France, in July 2016. The 9 revised full papers were carefully reviewed and selected from 50 submissions. The papers deal with the following topics: databases, data warehousing, data mining, data management, data security, knowledge and information systems and technologies; advanced application of data.

Database and information systems technologies have been rapidly evolving in several directions over the past years. New types and kinds of data, new types of applications and information systems to support them raise diverse challenges to be addressed. The so-called big data challenge, streaming data management and processing, social networks and other complex data analysis, including semantic reasoning into information systems supporting for instance trading, negotiations, and bidding mechanisms are just some of the emerging research topics. This volume contains papers contributed by six workshops: ADBIS Workshop on GPUs in Databases (GID 2012), Mining Complex and Stream Data (MCS'D'12),

Get Free Accelerating Sql Database Operations On A Gpu With Cuda

International Workshop on Ontologies meet Advanced Information Systems (OAIS'2012), Second Workshop on Modeling Multi-commodity Trade: Data models and processing (MMT'12), 1st ADBIS Workshop on Social Data Processing (SDP'12), 1st ADBIS Workshop on Social and Algorithmic Issues in Business Support (SAIBS), and the Ph.D. Consortium associated with the ADBIS 2012 conference that report on the recent developments and an ongoing research in the aforementioned areas.

?This two-volume set (CCIS 1075 and CCIS 1076) constitutes the refereed proceedings of the Third International Conference on Advanced Informatics for Computing Research, ICAICR 2019, held in Shimla, India, in June 2019. The 78 revised full papers presented were carefully reviewed and selected from 382 submissions. The papers are organized in topical sections on computing methodologies; hardware; information systems; networks; software and its engineering.

The main objective of the ICITMS 2012 is to provide a platform for researchers, engineers, academics and industrial professionals from all over the world to present their research results and development activities in Information Technology and Management Science. This conference provides opportunities for the delegates to exchange new ideas and application experiences face to face, to establish business or research relations and to find global partners for future collaboration.

The book includes the insights that reflect 'Advances in Computer and Computational Sciences' from upcoming researchers and leading academicians across the globe. It contains the high-quality peer-reviewed papers of 'International Conference on Computer, Communication and Computational Sciences (IC4S 2017), held during 11-12 October, 2017 in Thailand. These papers are arranged in the form of chapters. The content of this book is divided into two volumes that cover variety of topics such as intelligent hardware and software design, advanced communications, intelligent computing techniques, intelligent image processing, and web and informatics. This book helps the perspective readers' from computer industry and academia to derive the advances of next generation computer and communication technology and shape them into real life applications.

This book contains some selected papers from the International Conference on Extreme Learning Machine (ELM) 2017, held in Yantai, China, October 4-7, 2017. The book covers theories, algorithms and applications of ELM. Extreme Learning Machines (ELM) aims to enable pervasive learning and pervasive intelligence. As advocated by ELM theories, it is exciting to see the convergence of machine learning and biological learning from the long-term point of view. ELM may be one of the fundamental 'learning particles' filling the gaps between machine learning and biological learning (of which activation functions are even unknown). ELM represents a suite of (machine and biological) learning techniques in which hidden neurons need not be tuned: inherited from their ancestors or randomly generated. ELM learning theories show that effective learning algorithms can be derived based on randomly generated hidden neurons (biological neurons, artificial neurons, wavelets, Fourier series, etc) as long as they are nonlinear piecewise continuous, independent of training data and application environments. Increasingly, evidence from neuroscience suggests that similar principles apply in biological learning systems. ELM theories and algorithms argue that "random hidden neurons" capture an essential aspect of biological learning mechanisms as well as the intuitive sense that the efficiency of biological learning need not

Get Free Accelerating Sql Database Operations On A Gpu With Cuda

rely on computing power of neurons. ELM theories thus hint at possible reasons why the brain is more intelligent and effective than current computers. This conference will provide a forum for academics, researchers and engineers to share and exchange R&D experience on both theoretical studies and practical applications of the ELM technique and brain learning. It gives readers a glance of the most recent advances of ELM.

A Multi-Processor System-on-Chip (MPSoC) is the key component for complex applications. These applications put huge pressure on memory, communication devices and computing units. This book, presented in two volumes - Architectures and Applications - therefore celebrates the 20th anniversary of MPSoC, an interdisciplinary forum that focuses on multi-core and multi-processor hardware and software systems. It is this interdisciplinarity which has led to MPSoC bringing together experts in these fields from around the world, over the last two decades. Multi-Processor System-on-Chip 2 covers application-specific MPSoC design, including compilers and architecture exploration. This second volume describes optimization methods, tools to optimize and port specific applications on MPSoC architectures. Details on compilation, power consumption and wireless communication are also presented, as well as examples of modeling frameworks and CAD tools. Explanations of specific platforms for automotive and real-time computing are also included.

This book constitutes the refereed proceedings of the 12th International Conference entitled Beyond Databases, Architectures and Structures, BDAS 2016, held in Ustroń, Poland, in May/June 2016. It consists of 57 carefully reviewed papers selected from 152 submissions. The papers are organized in topical sections, namely artificial intelligence, data mining and knowledge discovery; architectures, structures and algorithms for efficient data processing; data warehousing and OLAP; natural language processing, ontologies and semantic Web; bioinformatics and biomedical data analysis; data processing tools; novel applications of database systems.

There is enormous pressure today for businesses across all industries to cut costs, enhance business performance, and deliver greater value with fewer resources. To take business analytics to the next level and drive tangible improvements to the bottom line, it is important to manage not only the volume of data, but the speed with which actionable findings can be drawn from a wide variety of disparate sources. The findings must be easily communicated to those responsible for making both strategic and tactical decisions. At the same time, strained IT budgets require that the solution be self-service for everyone from DBAs to business users, and easily deployed to thin, browser-based clients. Business analytics hosted in the Query Management Facility™ (QMFTM) on DB2® and System z® allow you to tackle these challenges in a practical way, using new features and functions that are easily deployed across the enterprise and easily consumed by business users who do not have prior IT experience. This IBM® Redbooks® publication provides step-by-step instructions on using these new features: Access to data that resides in any JDBC-compliant data source OLAP access through XMLA 150+ new analytical functions Graphical query interfaces and graphical reports Graphical, interactive dashboards Ability to integrate QMF functions with third-party applications Support for the IBM DB2 Analytics Accelerator A new QMF Classic perspective in QMF for Workstation Ability to start QMF for TSO as a DB2 for z/OS stored procedure New metadata capabilities, including ER diagrams and capability to federate data into a single virtual source

This volume LNCS 12925 constitutes the papers of the 23rd International Conference on Big Data Analytics and Knowledge

Get Free Accelerating Sql Database Operations On A Gpu With Cuda

Discovery, held in September 2021. Due to COVID-19 pandemic it was held virtually. The 12 full papers presented together with 15 short papers in this volume were carefully reviewed and selected from a total of 71 submissions. The papers reflect a wide range of topics in the field of data integration, data warehousing, data analytics, and recently big data analytics, in a broad sense. The main objectives of this event are to explore, disseminate, and exchange knowledge in these fields.

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

This book constitutes the proceedings of the 4th International Workshop on Motion in Games, held in Edinburgh, UK, in November 2011. The 30 revised full papers presented together with 8 revised poster papers in this volume were carefully reviewed and selected from numerous submissions. The papers are organized in topical sections on character animation, motion synthesis, physically-based character motion, behavior animation, animation systems, crowd simulation, as well as path planning and navigation.

Transforming data from operational data models to purpose-oriented data structures has been commonplace for the last decades. Data transformations are heavily used in all types of industries to provide information to various users at different levels. Depending on individual needs, the transformed data is stored in various different systems. Sending operational data to other systems for further processing is then required, and introduces much complexity to an existing information technology (IT) infrastructure. Although maintenance of additional hardware and software is one component, potential inconsistencies and individually managed refresh cycles are others. For decades, there was no simple and efficient way to perform data transformations on the source system of operational data. With IBM® DB2® Analytics

Get Free Accelerating Sql Database Operations On A Gpu With Cuda

Accelerator, DB2 for z/OS is now in a unique position to complete these transformations in an efficient and well-performing way. DB2 for z/OS completes these while connecting to the same platform as for operational transactions, helping you to minimize your efforts to manage existing IT infrastructure. Real-time analytics on incoming operational transactions is another demand. Creating a comprehensive scoring model to detect specific patterns inside your data can easily require multiple iterations and multiple hours to complete. By enabling a first set of analytical functionality in DB2 Analytics Accelerator, those dedicated mining algorithms can now be run on an accelerator to efficiently perform these modeling tasks. Given the speed of query processing on an accelerator, these modeling tasks can now be performed much quicker compared to traditional relational database management systems. This speed enables you to keep your scoring algorithms more up-to-date, and ultimately adapt more quickly to constantly changing customer behaviors. This IBM Redbooks® publication describes the new table type that is introduced with DB2 Analytics Accelerator V4.1 PTF5 that enables more efficient data transformations. These tables are called accelerator-only tables, and can exist on an accelerator only. The tables benefit from the accelerator performance characteristics, while maintaining access through existing DB2 for z/OS application programming interfaces (APIs). Additionally, we describe the newly introduced analytical capabilities with DB2 Analytics Accelerator V5.1, putting you in the position to efficiently perform data modeling for online analytical requirements in your DB2 for z/OS environment. This book is intended for technical decision-makers who want to get a broad understanding about the analytical capabilities and accelerator-only tables of DB2 Analytics Accelerator. In addition, you learn about how these capabilities can be used to accelerate in-database transformations and in-database analytics in various environments and scenarios, including the following scenarios: Multi-step processing and reporting in IBM DB2 Query Management Facility™, IBM Campaign, or Microstrategy environments In-database transformations using IBM InfoSphere® DataStage® Ad hoc data analysis for data scientists In-database analytics using IBM SPSS® Modeler

This book constitutes the proceedings of the Third Latin American Conference on High Performance Computing, CARLA 2016, held in Mexico City, Mexico, in August/September 2016. The 30 papers presented in this volume were carefully reviewed and selected from 70 submissions. They are organized in topical sections named: HPC Infrastructure and Applications; Parallel Algorithms and Applications; HPC Applications and Simulations.

The book gathers papers addressing state-of-the-art research in all areas of Information and Communication Technologies and their applications in intelligent computing, cloud storage, data mining and software analysis. It presents the outcomes of the third International Conference on Information and Communication Technology for Intelligent Systems, which was held on April 6-7, 2018, in Ahmedabad, India. Divided into two volumes, the book discusses the fundamentals of various data analytics and algorithms, making it a valuable resource for researchers' future studies.

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

Get Free Accelerating Sql Database Operations On A Gpu With Cuda

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 special issue contains extended and revised versions of 4 papers, selected from the 25 papers presented at the satellite events associated with the 17th East-European Conference on Advances in Databases and Information Systems (ADBIS 2013), held on September 1-4, 2013 in Genoa, Italy. The three satellite events were GID 2013, the Second International Workshop on GPUs in Databases; SoBI 2013, the First International Workshop on Social Business Intelligence: Integrating Social Content in Decision Making; and OAIS 2013, the Second International Workshop on Ontologies Meet Advanced Information Systems. The papers cover various topics in large-scale data and knowledge-centered systems, including GPU-accelerated database systems and GPU-based compression for large time series databases, design of parallel data warehouses, and schema matching. The special issue content, which combines both theoretical and application-based contributions, gives a useful overview of some of the current trends in large-scale data and knowledge management and will stimulate new ideas for further research and development within both the scientific and industrial communities.

An expert guide for IT administrators needing to create and manage a public cloud and virtual network using Microsoft Azure With Microsoft Azure challenging Amazon Web Services (AWS) for market share, there has been no better time for IT professionals to broaden and expand their knowledge of Microsoft's flagship virtualization and cloud computing service. Microsoft Azure Infrastructure Services for Architects: Designing Cloud Solutions helps readers develop the skills required to understand the capabilities of Microsoft Azure for Infrastructure Services and implement a public cloud to achieve full virtualization of data, both on and off premise. Microsoft Azure provides granular control in choosing core infrastructure components, enabling IT administrators to deploy new Windows Server and Linux virtual machines, adjust usage as requirements change, and scale to meet the infrastructure needs of their entire organization. This accurate, authoritative book covers topics including IaaS cost and options, customizing VM storage, enabling external connectivity to Azure virtual machines, extending Azure Active Directory, replicating and backing up to Azure, disaster recovery, and much more. New users and experienced professionals alike will: Get expert guidance on understanding, evaluating, deploying, and maintaining Microsoft Azure environments from Microsoft MVP and technical specialist John Savill Develop the skills to set up cloud-based virtual machines, deploy web servers, configure hosted data stores, and use other key Azure technologies Understand how to design and implement serverless and hybrid solutions Learn to use enterprise security guidelines for Azure deployment Offering the most up to date information and practical advice, Microsoft Azure Infrastructure Services for Architects: Designing Cloud Solutions is an essential resource for IT administrators, consultants and engineers responsible for learning, designing, implementing, managing, and maintaining Microsoft virtualization and cloud technologies.

This two-volume set, LNCS 9658 and 9659, constitutes the thoroughly refereed proceedings of the 17th International Conference on Web-Age Information Management, WAIM 2016, held in Nanchang, China, in June 2016. The 80 full research papers presented together with 8 demonstrations were carefully reviewed and selected from 266 submissions. The focus of the conference is on following topics: data mining, spatial and temporal databases, recommender systems, graph data

Get Free Accelerating Sql Database Operations On A Gpu With Cuda

management, information retrieval, privacy and trust, query processing and optimization, social media, big data analytics, and distributed and cloud computing.

This book constitutes the revised selected papers of the 12th International Workshop on Information Search, Integration and Personalization, ISIP 2018, held in Fukuoka, Japan, in May 2018. The volume presents 1 invited paper as well as 7 revised full papers, which were carefully reviewed and selected from 13 papers submitted to these post-conference proceedings. The papers are organized in topical sections on data integration; text and document management; advanced data mining techniques.

Expert Oracle RAC 12c is a hands-on book helping you understand and implement Oracle Real Application Clusters (RAC), and to reduce the total-cost-of-ownership (TCO) of a RAC database. As a seasoned professional, you are probably aware of the importance of understanding the technical details behind the RAC stack. This book provides deep understanding of RAC concepts and implementation details that you can apply toward your day-to-day operational practices. You'll be guided in troubleshooting and avoiding trouble in your installation. Successful RAC operation hinges upon a fast-performing network interconnect, and this book dedicates a chapter solely to that very important and easily overlooked topic. All four authors are experienced RAC engineers with a wealth of hard-won experience encountering and surmounting the challenges of running a RAC environment that delivers on its promise. In Expert Oracle RAC 12c they provide you a framework in which to avoid repeating their hard-won lessons. Their goal is for you to manage your own RAC environment with ease and expertise. Provides a deep conceptual understanding of RAC Provides best practices to implement RAC properly and match application workload Enables readers to troubleshoot RAC with ease What you'll learn Know when to apply RAC, and when not to Design applications to take advantage of RAC Troubleshoot and solve clusterware problems Manage database backup and recovery in RAC Stay on top of locking issues and deadlock detection Harness the performance from parallel processing in RAC Support your RAC environment with a healthy network interconnect Who this book is for Expert Oracle RAC 12c is for experienced Oracle Database Administrators (DBAs) who are ready to take the next step in their career by expanding their skill set to include building and managing Oracle Real Application Clusters (RAC). DBAs and architects who are in the process of implementing RAC can immensely benefit from this book. It's an excellent choice for DBAs to learn RAC conceptually, understand best practices, and become experts in troubleshooting RAC problems. Table of Contents Overview of Oracle RAC Clusterware Management and Troubleshooting RAC Operational Practices RAC New Features Storage and ASM Practices Application Design Issues Managing and Optimizing a Complex RAC Environment Backup and Recovery in RAC Network Practices in RAC RAC Database Optimization Locks and Deadlocks Parallel Query in RAC Clusterware and Database Upgrades Oracle RAC One Node

This proceedings set contains selected Computer, Information and Education Technology related papers from the 2014 International Conference on Computer, Intelligent Computing and Education Technology (CICET 2014), held March 27-28, 2014 in Hong Kong. The proceedings aims to provide a platform for researchers, engineers and academics as well as indu