

## Database Comparison Dell Poweredge R710 Vs Hp Proliant

The LNCS 12115 constitutes the workshop papers which were held also online in conjunction with the 25th International Conference on Database Systems for Advanced Applications in September 2020. The complete conference includes 119 full papers presented together with 19 short papers plus 15 demo papers and 4 industrial papers in this volume were carefully reviewed and selected from a total of 487 submissions. DASFAA 2020 presents this year following five workshops: The 7th International Workshop on Big Data Management and Service (BDMS 2020) The 6th International Symposium on Semantic Computing and Personalization (SeCoP 2020) The 5th Big Data Quality Management (BDQM 2020) The 4th International Workshop on Graph Data Management and Analysis (GDMA 2020) The 1st International Workshop on Artificial Intelligence for Data Engineering (AIDE 2020)

This book constitutes the workshop proceedings of the 21st International Conference on Database Systems for Advanced Applications, DASFAA 2016, held in Dallas, TX, USA, in April 2016. The volume contains 32 full papers (selected from 43 submissions) from 4 workshops, each focusing on a specific area that contributes to the main themes of DASFAA 2016: The Third International Workshop on Semantic Computing and Personalization, SeCoP 2016; the Third International Workshop on Big Data Management and Service, BDMS 2016; the First International Workshop on Big Data Quality Management, BDQM 2016; and the Second International Workshop on Mobile of Internet, Mol 2016.

This book focuses on the core question of the necessary architectural support provided by hardware to efficiently run virtual machines, and of the corresponding design of the hypervisors that run them. Virtualization is still possible when the instruction set architecture lacks such support, but the hypervisor remains more complex and must rely on additional techniques. Despite the focus on architectural support in current architectures, some historical perspective is necessary to appropriately frame the problem. The first half of the book provides the historical perspective of the theoretical framework developed four decades ago by Popek and Goldberg. It also describes earlier systems that enabled virtualization despite the lack of architectural support in hardware. As is often the case, theory defines a necessary—but not sufficient—set of features, and modern architectures are the result of the combination of the theoretical framework with insights derived from practical systems. The second half of the book describes state-of-the-art support for virtualization in both x86-64 and ARM processors. This book includes an in-depth description of the CPU, memory, and I/O virtualization of these two processor architectures, as well as case studies on the Linux/KVM, VMware, and Xen hypervisors. It concludes with a performance comparison of virtualization on current-generation x86- and ARM-based systems across multiple hypervisors.

Summary PowerShell and WMI is an example-driven guide for administrators managing networks of Windows servers and desktops. With 150 practical examples, including ready-to-reuse scripts and techniques, you'll learn the ins and outs of automating WMI via PowerShell v3. You'll also find deep coverage of all aspects of Windows administration, including IIS, DNS and Hyper-V. About the Technology WMI, on its own, is simply a collection of Windows management facilities. Paired with PowerShell, however, WMI becomes a brilliant toolset for automating servers, networks, and remote Windows computers. About the Book PowerShell and WMI is an industrial-strength guide for administrators of Windows networks, servers, and desktops. You'll start with practical overviews of PowerShell and of WMI. Then you'll explore 150 specific examples—all with ready-to-use scripts—designed to simplify your day-to-day system management. Each tested technique is configured to load as part of a PowerShell module. A set of handy appendixes includes references for PowerShell and WMI. Prior exposure to PowerShell and WMI is helpful but not required. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. What's Inside Managing Windows, IIS, and Hyper-V Remote desktops and devices Logs, jobs, and performance New PowerShell v3 WMI functionality ===== Table of Contents PART 1 TOOLS OF THE TRADE Solving administrative challenges Using PowerShell WMI in depth Best practices and optimization PART 2 WMI IN THE ENTERPRISE? System documentation Disk systems Registry administration Filesystem administration Services and processes Printers Configuring network adapters Managing IIS Configuring a server Users and security Logs, jobs, and performance Administering Hyper-V with PowerShell and WMI PART 3 THE FUTURE: POWERSHELL V3 AND WMI WMI over WSMAN Your own WMI cmdlets CIM cmdlets and sessions

The LNCS series reports state-of-the-art results in computer science research, development, and education, at a high level and in both printed and electronic form. Enjoying tight cooperation with the R&D community, with numerous individuals, as well as with prestigious organizations and societies, LNCS has grown into the most comprehensive computer science research forum available. The scope of LNCS, including its subseries LNAI and LNBI, spans the whole range of computer science and information technology including interdisciplinary topics in a variety of application fields. The type of material published traditionally includes More recently, several color-cover sublines have been added featuring, beyond a collection of papers, various added-value components; these sublines include In parallel to the printed book, each new volume is published electronically in LNCS Online.

vSphere High Performance Cookbook is written in a practical, helpful style with numerous recipes focusing on answering and providing solutions to common, and not-so common, performance issues and problems. The book is primarily written for technical professionals with system administration skills and some VMware experience who wish to learn about advanced optimization and the configuration features and functions for vSphere 5.1.

A guide for MySQL administrators covers such topics as benchmarking, server performance, indexing, queries, hardware optimization, replication, scaling, cloud hosting, and backup and recovery.

Bring on-premise and cloud collaboration features to life with Microsoft's enterprise content management platform - SharePoint Server Key Features Get up to speed with the latest version of SharePoint and make the most of its features Set up and effectively manage your SharePoint Farm in the cloud or on premise Learn how to align your development tools and cloud infrastructure to support collaboration Book Description Microsoft's latest addition to their product range, SharePoint Server 2019, is a new enterprise content management platform that brings on-premise collaboration features to life. It can be used as an isolated platform or in a hybrid connected configuration providing management and connectivity to Office 365. You can use the SharePoint framework to host sites, information, data, and applications in a robust CMS that centralizes collaborative content for enterprises. SharePoint 2019 enables new integrations and features that will allow you to work seamlessly with new and old Office products such as Microsoft Power Apps and other Microsoft Office applications. Implementing Microsoft SharePoint 2019 will help you understand the challenges, planning, migration steps, installation concepts, and configuration involved in providing this platform for your enterprise. The book will also show you what the platform brings to the table from an on-premise server perspective. If you're new to SharePoint 2019, you'll also be guided through how to get servers up and running so that you and your user community can become

productive with this powerful new platform. By the end of this book, you'll be well-versed in Microsoft SharePoint 2019 and have the knowledge you need to apply your skills in the real world. What you will learn Understand changes to the platform and how to migrate from other versions of SharePoint Explore infrastructure planning and governance relating to collaborative environments Install and configure network components, servers, and desktops Use SharePoint services and other Microsoft product servers and apps Monitor and troubleshoot SharePoint after it is implemented Discover the tools that can be used with SharePoint 2019 for BI and reporting Delve into social features and collaboration Maintain, monitor, and support the rollout of the platform in your enterprise Who this book is for The book is for SharePoint administrators, project managers, and architects who have some experience in designing, planning, implementing, and managing SharePoint Farms.

This book constitutes the refereed proceedings of the 4th Annual International Symposium on Information Management and Big Data, SIMBig 2017, held in Lima, Peru, in September 2017. The 10 revised full papers presented were carefully reviewed and selected from 71 submissions. The papers address issues such as Data Science, Big Data, Data Mining, Natural Language Processing, Text Mining, Information Retrieval, Machine Learning, Semantic Web, Ontologies, Web Mining, Knowledge Representation and Linked Open Data, Social Web and Web Science, Information Visualization.

The de novo fabrication of custom DNA molecules is a transformative technology that significantly affects the biotechnology industry. Basic genetic engineering techniques for manipulating DNA in vitro opened an incredible field of opportunity in the life sciences. In, Gene Synthesis: Methods and Protocols expert researchers in the field detail many of the methods which are now commonly used to fabricate DNA . These include methods and techniques for the assembly of oligonucleotide, cloning of synthons into larger fragments, protocols and software applications, and educational and biosecurity impacts of gene synthesis. Written in the highly successful Methods in Molecular Biology™ series format, the chapters include the kind of detailed description and implementation advice that is crucial for getting optimal results in the laboratory. Thorough and intuitive, Gene Synthese: Methods and Protocols aids scientists in understanding all the different stages of a complex gene synthesis process, while refining their understanding of gene synthesis and determine what part of the process they can or should do in their laboratory and what parts should be contracted to a specialized service provider.

Database Systems for Advanced Applications DASFAA 2016 International Workshops: BDMS, BDQM, MoI, and SeCoP, Dallas, TX, USA, April 16-19, 2016, Proceedings Springer A practical guide packed with step-by-step recipes to design a virtual datacenter using VMware 5.x. This book is a guide for anyone interested in designing virtualized datacenters using VMware vSphere 5.x and the supporting components. Current administrators of VMware vSphere environments will find this book useful when interested in becoming a vSphere Architect or are interested in learning more about the virtual datacenter design process. Knowledge of vSphere installation, configuration, and administration is a prerequisite.

Over the last ten years, the ARM architecture has become one of the most pervasive architectures in the world, with more than 2 billion ARM-based processors embedded in products ranging from cell phones to automotive braking systems. A world-wide community of ARM developers in semiconductor and product design companies includes software developers, system designers and hardware engineers. To date no book has directly addressed their need to develop the system and software for an ARM-based system. This text fills that gap. This book provides a comprehensive description of the operation of the ARM core from a developer's perspective with a clear emphasis on software. It demonstrates not only how to write efficient ARM software in C and assembly but also how to optimize code. Example code throughout the book can be integrated into commercial products or used as templates to enable quick creation of productive software. The book covers both the ARM and Thumb instruction sets, covers Intel's XScale Processors, outlines distinctions among the versions of the ARM architecture, demonstrates how to implement DSP algorithms, explains exception and interrupt handling, describes the cache technologies that surround the ARM cores as well as the most efficient memory management techniques. A final chapter looks forward to the future of the ARM architecture considering ARMv6, the latest change to the instruction set, which has been designed to improve the DSP and media processing capabilities of the architecture. \* No other book describes the ARM core from a system and software perspective. \* Author team combines extensive ARM software engineering experience with an in-depth knowledge of ARM developer needs. \* Practical, executable code is fully explained in the book and available on the publisher's Website. \* Includes a simple embedded operating system.

VMware ESX Server in the Enterprise Planning and Securing Virtualization Servers The Most Complete, Practical, Solutions-Focused Guide to Running ESX Server 3 VMware ESX Server in the Enterprise is the definitive, real-world guide to planning, deploying, and managing today's leading virtual infrastructure platform in mission-critical environments. Drawing on his extensive experience consulting on large-scale ESX Server implementations, Edward L. Haletky brings together an unprecedented collection of tips, best practices, and field-tested solutions. More than any other author, he illuminates the real issues, tradeoffs, and pitfalls associated with ESX Server—and shows how to make the most of it in your unique environment. Haletky covers the entire lifecycle: planning, installation, system monitoring, tuning, clustering, security, disaster recovery, and much more. Throughout, he supports his recommendations with examples from real-world deployments. He also provides detailed checklists for handling crucial issues such as caching, networking, storage, and hardware selection. Many of his techniques and practices apply to all current virtualization platforms, not just ESX Server. This book will be an indispensable resource for every network architect, administrator, and IT professional who works with virtual servers. ESX Server newcomers will find the soup-to-nuts introduction they desperately need; experienced users will find an unparalleled source of field-tested answers and solutions. In this book, you'll learn how to: • Identify key

differences between ESX v3.x.y and ESX v2.5.x and their implications • Perform a complete installation—with automated scripting techniques and samples • Efficiently audit, monitor, and secure ESX Server • Discover SAN storage pitfalls and solutions—with detailed guidance for specific SANs, switches, and fibre-channel adapters • Understand ESX Server networking: NIC teaming, vSwitches, network lag, and troubleshooting • Configure ESX Server via the Management User Interface, Virtual Center client, and command line interface • Install Windows, Linux, and NetWare VMs: prepare media images, place configuration files, handle sizing and swap files, and more • Use Dynamic Resource Load Balancing to consistently achieve utilization goals • Implement effective backup and disaster recovery procedures Edward L. Haletky owns AstroArch Consulting, Inc., a consultancy specializing in virtualization, security, and networking. He has been rated by his peers on the VMware Discussion Forums as a “virtuoso” for his work in answering VMware security and configuration questions. Prior to establishing AstroArch, Haletky was a member of Hewlett-Packard’s Virtualization, Linux, and High-Performance Technical Computing teams. He holds a degree in Aeronautical and Astronautical Engineering from Purdue University.

Hands-on troubleshooting methods on the most recent release of SQL Server The 2012 release of SQL Server is the most significant one since 2005 and introduces an abundance of new features. This critical book provides in-depth coverage of best practices for troubleshooting performance problems based on a solid understanding of both SQL Server and Windows internals and shows experienced DBAs how to ensure reliable performance. The team of authors shows you how to master the use of specific troubleshooting tools and how to interpret their output so you can quickly identify and resolve any performance issue on any server running SQL Server. Covers the core technical topics required to understand how SQL Server and Windows should be working Shares best practices so that you know how to proactively monitor and avoid problems Shows how to use tools to quickly gather, analyze, and effectively respond to the source of a system-wide performance issue Professional SQL Server 2012 Internals and Troubleshooting helps you to quickly become familiar with the changes of this new release so that you can best handle database performance and troubleshooting.

Computer Architecture: A Quantitative Approach, Sixth Edition has been considered essential reading by instructors, students and practitioners of computer design for over 20 years. The sixth edition of this classic textbook from Hennessy and Patterson, winners of the 2017 ACM A.M. Turing Award recognizing contributions of lasting and major technical importance to the computing field, is fully revised with the latest developments in processor and system architecture. The text now features examples from the RISC-V (RISC Five) instruction set architecture, a modern RISC instruction set developed and designed to be a free and openly adoptable standard. It also includes a new chapter on domain-specific architectures and an updated chapter on warehouse-scale computing that features the first public information on Google's newest WSC. True to its original mission of demystifying computer architecture, this edition continues the longstanding tradition of focusing on areas where the most exciting computing innovation is happening, while always keeping an emphasis on good engineering design. Winner of a 2019 Textbook Excellence Award (Texty) from the Textbook and Academic Authors Association Includes a new chapter on domain-specific architectures, explaining how they are the only path forward for improved performance and energy efficiency given the end of Moore’s Law and Dennard scaling Features the first publication of several DSAs from industry Features extensive updates to the chapter on warehouse-scale computing, with the first public information on the newest Google WSC Offers updates to other chapters including new material dealing with the use of stacked DRAM; data on the performance of new NVIDIA Pascal GPU vs. new AVX-512 Intel Skylake CPU; and extensive additions to content covering multicore architecture and organization Includes "Putting It All Together" sections near the end of every chapter, providing real-world technology examples that demonstrate the principles covered in each chapter Includes review appendices in the printed text and additional reference appendices available online Includes updated and improved case studies and exercises ACM named John L. Hennessy and David A. Patterson, recipients of the 2017 ACM A.M. Turing Award for pioneering a systematic, quantitative approach to the design and evaluation of computer architectures with enduring impact on the microprocessor industry

"This Technical Manual is a guide to personnel who train post fire fighters. It consists of a series of lesson plans covering each phase of firemanship, and a brief explanation of how to use the plans."--Introd. This book constitutes the refereed proceedings of the 11th International Conference on Similarity Search and Applications, SISAP 2018, held in Lima, Peru, in October 2018. The 16 full papers presented together with 3 short papers and 1 demonstration paper were carefully reviewed and selected from 31 submissions. The papers deal with issues surrounding the theory, design, analysis, practice, and application of content-based and feature-based similarity search. They are organized in the following topical sections: metric search; visual search; nearest neighbor queries; clustering and outlier detection; graphs and applications; and shared session SISAP and SPIRE.

Mastering Ubuntu Server, Third Edition not only strengthens your server fundamentals but also equips you with the advanced concepts of Ubuntu 20.04 LTS. It polishes and expands your skill set to prepare you for better business opportunities.

This book is dedicated to advances in the field of user authentication. The book covers detailed description of the authentication process as well as types of authentication modalities along with their several features (authentication factors). It discusses the use of these modalities in a time-varying operating environment, including factors such as devices, media and surrounding conditions, like light, noise, etc. The book is divided into several parts that cover descriptions of several biometric and non-biometric authentication modalities, single factor and multi-factor authentication systems (mainly, adaptive), negative authentication system, etc. Adaptive strategy ensures the incorporation of the existing environmental conditions on the selection of authentication factors and provides significant diversity in the selection process. The contents of this book will prove useful to practitioners, researchers and students. The book is suited to be used a text in advanced/graduate courses on User Authentication Modalities. It can also be used as a textbook for professional development and certification coursework for practicing engineers and computer scientists.

Get up to date with the finer points of Ubuntu Server using this comprehensive guide About This Book Get well-versed with newly-added features in Ubuntu 16.04 Master the art of installing, managing, and troubleshooting Ubuntu Server A practical easy-to-understand book that will help you enhance your existing skills. Who This Book Is For This book is intended for readers with intermediate or advanced-beginner skills with Linux, who would like to learn all about setting up servers with Ubuntu Server. This book assumes that the reader knows the basics of Linux, such as editing configuration files and running basic commands. What You Will Learn Learn how to manage users, groups, and permissions Encrypt and decrypt disks with Linux Unified Key Setup /Luks Setup SSH for remote access, and connect it to other nodes Understand how to add, remove, and search for packages Use NFS and Samba to share directories with other users Get to know techniques for managing Apache and MariaDB Explore best practices and troubleshooting techniques In Detail Ubuntu is a Debian-based Linux operating system, and has various versions targeted at servers, desktops, phones, tablets and televisions. The Ubuntu Server Edition, also called Ubuntu Server, offers support for several common configurations, and also simplifies common Linux server deployment processes. With this book as their guide, readers will be able to configure and deploy Ubuntu Servers using Ubuntu Server 16.04, with all the skills necessary to manage real servers. The book begins with the concept of user management, group management, as well as file-system permissions. To manage your storage on Ubuntu Server systems, you will learn how to add and format storage and view disk usage. Later, you will also learn how to configure network

interfaces, manage IP addresses, deploy Network Manager in order to connect to networks, and manage network interfaces. Furthermore, you will understand how to start and stop services so that you can manage running processes on Linux servers. The book will then demonstrate how to access and share files to or from Ubuntu Servers. You will learn how to create and manage databases using MariaDB and share web content with Apache. To virtualize hosts and applications, you will be shown how to set up KVM/Qemu and Docker and manage virtual machines with virt-manager. Lastly, you will explore best practices and troubleshooting techniques when working with Ubuntu Servers. By the end of the book, you will be an expert Ubuntu Server user well-versed in its advanced concepts. Style and Approach This book is an advanced guide that will show readers how to administer, manage, and deploy Ubuntu server and will also provide expert-level knowledge on advanced security and backup techniques.

The two-volume set, LNCS 9326 and LNCS 9327 constitutes the refereed proceedings of the 20th European Symposium on Research in Computer Security, ESORICS 2015, held in Vienna, Austria, in September 2015. The 59 revised full papers presented were carefully reviewed and selected from 298 submissions. The papers address issues such as networks and Web security; system security; crypto application and attacks; risk analysis; privacy; cloud security; protocols and attribute-based encryption; code analysis and side-channels; detection and monitoring; authentication; policies; and applied security.

The computing world today is in the middle of a revolution: mobile clients and cloud computing have emerged as the dominant paradigms driving programming and hardware innovation today. The Fifth Edition of Computer Architecture focuses on this dramatic shift, exploring the ways in which software and technology in the cloud are accessed by cell phones, tablets, laptops, and other mobile computing devices. Each chapter includes two real-world examples, one mobile and one datacenter, to illustrate this revolutionary change. Updated to cover the mobile computing revolution Emphasizes the two most important topics in architecture today: memory hierarchy and parallelism in all its forms. Develops common themes throughout each chapter: power, performance, cost, dependability, protection, programming models, and emerging trends ("What's Next") Includes three review appendices in the printed text. Additional reference appendices are available online. Includes updated Case Studies and completely new exercises.

Ubuntu is a Debian-based Linux distribution with versions available for both desktops as well as servers. The Server edition, Ubuntu Server, has set the industry standard for Linux in the data center as well as the cloud. Organizations, inventors, and hobbyists alike will benefit from its flexible configuration, fast deployment, and a plethora ...

UNESCO issued this publication to demystify the concept of open access (OA) and to provide concrete steps on putting relevant policies in place. Its focus is on scientific research from peer-reviewed journal articles. Building capacities in Member States for Open Access is a necessary but not sufficient condition for promotion of the concept. Creating an enabling policy environment for OA is therefore a priority. This publication will serve the needs of OA policy development at the government, institutional and funding agency level. The overall objective of the Policy Guidelines is to promote Open Access in Member States by facilitating understanding of all relevant issues related to Open Access. The guidelines are not prescriptive in nature, but are suggestive to facilitate knowledge-based decision-making to adopt OA policies and strengthen national research systems.

This book constitutes the refereed proceedings of three workshops held at the 19th International Conference on Financial Cryptography and Data Security, FC 2015, in San Juan, Puerto Rico, in January 2015. The 22 full papers presented were carefully reviewed and selected from 39 submissions. They feature the outcome of the Second Workshop on Bitcoin Research, BITCOIN 2015, the Third Workshop on Encrypted Computing and Applied Homomorphic Cryptography, WAHC 2015, and the First Workshop on Wearable Security and Privacy, Wearable 2015.

This book constitutes the refereed proceedings of the 29th Annual IFIP WG 11.3 International Working Conference on Data and Applications Security and Privacy, DBSec 2015, held in Fairfax, VA, USA, in July 2015. The 18 revised full papers and 6 short papers presented were carefully reviewed and selected from 45 submissions. The papers are organized in the following topical sections: data anonymization and computation; access control and authorization; user privacy; authentication and information integration; privacy and trust; access control and usage policies; network and internet security; and information flow and inference.

Learn the art of efficient web scraping and crawling with Python About This Book Extract data from any source to perform real time analytics. Full of techniques and examples to help you crawl websites and extract data within hours. A hands-on guide to web scraping and crawling with real-life problems and solutions Who This Book Is For If you are a software developer, data scientist, NLP or machine-learning enthusiast or just need to migrate your company's wiki from a legacy platform, then this book is for you. It is perfect for someone , who needs instant access to large amounts of semi-structured data effortlessly. What You Will Learn Understand HTML pages and write XPath to extract the data you need Write Scrapy spiders with simple Python and do web crawls Push your data into any database, search engine or analytics system Configure your spider to download files, images and use proxies Create efficient pipelines that shape data in precisely the form you want Use Twisted Asynchronous API to process hundreds of items concurrently Make your crawler super-fast by learning how to tune Scrapy's performance Perform large scale distributed crawls with scrapyd and scrapinghub In Detail This book covers the long awaited Scrapy v 1.0 that empowers you to extract useful data from virtually any source with very little effort. It starts off by explaining the fundamentals of Scrapy framework, followed by a thorough description of how to extract data from any source, clean it up, shape it as per your requirement using Python and 3rd party APIs. Next you will be familiarised with the process of storing the scrapped data in databases as well as search engines and performing real time analytics on them with Spark Streaming. By the end of this book, you will perfect the art of scarping data for your applications with ease Style and approach It is a hands on guide, with first few chapters written as a tutorial, aiming to motivate you and get you started quickly. As the book progresses, more advanced features are explained with real world examples that can be referred while developing your own web applications.

Aerospace propulsion devices embody some of the most advanced technologies, ranging from materials, fluid control, and heattransfer and combustion. In order to maximize the performance, sophisticated testing and computer simulation tools are developed and used. Aerospace Propulsion comprehensively covers the mechanics and thermal-fluid aspects of aerospace propulsion, starting from the fundamental principles, and covering applications to gas-turbine and space propulsion (rocket) systems. It presents modern analytical methods using MATLAB and other advanced software and includes essential elements of both gas-turbine and rocket propulsion systems. Gas turbine coverage includes thermodynamic analysis, turbine components, diffusers, compressors, turbines, nozzles, compressor-turbine matching, combustors and afterburners. Rocket coverage includes chemical rockets, electrical rockets, nuclear and solar sail. Key features: Both gas-turbine and rocket propulsion covered in a single volume Presents modern analytical methods and examples Combines fundamentals and applications, including space applications Accompanied by a website containing MATLAB examples, problem sets and solutions Aerospace Propulsion is a comprehensive textbook for senior undergraduate graduate and aerospace propulsion courses, and is also an excellent reference for researchers and

practicingengineers working in this area.

Balances old and new methods of chemical analysis by treating classic topics such as volumetric and gravimetric methods as well as newer areas including solvent extraction and chromatographic methods of separation. Emphasizes fundamental principles of each method and indicates possible applications to other areas of chemistry. It can be used as both a textbook for postgraduate students majoring in analytical chemistry and a reference for practicing analytical chemists and researchers.

The world of scholarship is changing rapidly. Increasing demands on scholars, the growing size and complexity of questions and problems to be addressed, and advances in sophistication of data collection, analysis, and presentation require new approaches to scholarship. A ubiquitous, open information infrastructure for scholarship, consisting of linked open data, open-source software tools, and a community committed to sustainability are emerging to meet the needs of scholars today. This book provides an introduction to VIVO, <http://vivoweb.org/>, a tool for representing information about research and researchers -- their scholarly works, research interests, and organizational relationships. VIVO provides an expressive ontology, tools for managing the ontology, and a platform for using the ontology to create and manage linked open data for scholarship and discovery. Begun as a project at Cornell and further developed by an NIH funded consortium, VIVO is now being established as an open-source project with community participation from around the world. By the end of 2012, over 20 countries and 50 organizations will provide information in VIVO format on more than one million researchers and research staff, including publications, research resources, events, funding, courses taught, and other scholarly activity. The rapid growth of VIVO and of VIVO-compatible data sources speaks to the fundamental need to transform scholarship for the 21st century. Table of Contents: Scholarly Networking Needs and Desires / The VIVO Ontology / Implementing VIVO and Filling It with Life / Case Study: University of Colorado at Boulder / Case Study: Weill Cornell Medical College / Extending VIVO / Analyzing and Visualizing VIVO Data / The Future of VIVO: Growing the Community

computationalmodelswith experimentaldata. A completedatasetwasprovided in advance, containing the cardiac geometry and ?bre orientations from MRI as well as epicardial transmembrane potentials from optical mapping.

SQL Server Hardware will provide the fundamental knowledge and resources you need to make intelligent decisions about choice, and optimal installation and configuration, of SQL Server hardware, operating system and the SQL Server RDBMS.

Foreword -- Foreword to the First Printing -- Preface -- Chapter 1 -- Introduction -- Chapter 2 -- Message Switching Layer -- Chapter 3 -- Deadlock, Livelock, and Starvation -- Chapter 4 -- Routing Algorithms -- Chapter 5 -- CollectiveCommunicationSupport -- Chapter 6 -- Fault-Tolerant Routing -- Chapter 7 -- Network Architectures -- Chapter 8 -- Messaging Layer Software -- Chapter 9 -- Performance Evaluation -- Appendix A -- Formal Definitions for Deadlock Avoidance -- Appendix B -- Acronyms -- References -- Index.

How can you bring out MySQL's full power? With High Performance MySQL, you'll learn advanced techniques for everything from designing schemas, indexes, and queries to tuning your MySQL server, operating system, and hardware to their fullest potential. This guide also teaches you safe and practical ways to scale applications through replication, load balancing, high availability, and failover. Updated to reflect recent advances in MySQL and InnoDB performance, features, and tools, this third edition not only offers specific examples of how MySQL works, it also teaches you why this system works a.

[Copyright: f36f669e922bcf8084e6071c2115fd76](http://www.mysql.com/doc/mysql-high-performance-guide/en/)