

Isilon Administration Guide

Right here, we have countless book Isilon Administration Guide and collections to check out. We additionally give variant types and as a consequence type of the books to browse. The up to standard book, fiction, history, novel, scientific research, as well as various supplementary sorts of books are readily affable here.

As this Isilon Administration Guide, it ends in the works bodily one of the favored book Isilon Administration Guide collections that we have. This is why you remain in the best website to look the unbelievable book to have.

Managing the Mail United States. National Archives and Records Service. Office of Records Management 1971

Dr. Tom Shinder's Configuring ISA Server 2004 Debra Littlejohn Shinder 2004-12-31 Dr. Tom and Debra Shinder have become synonymous with Microsoft's flagship firewall product ISA Server, as a result of Tom's prominent role as a member of the beta

development team, and Tom and Deb's featured placement on both Microsoft's ISA Server Web site and ISAserver.org. Tom and Deb's book on the first release of the product "Configuring ISA Server 2000" dominated the ISA Server 2000 book market having sold over 40,000 copies worldwide, and the ISA Server community is eagerly awaiting Tom and Deb's book on ISA Server 2004, which is the dramatically upgraded new release from Microsoft. Dr. Tom and Debra Shinder have become synonymous with Microsoft's flagship firewall product ISA Server, as a result of Tom's prominent role as a member of the beta development team, and Tom and Deb's featured placement on both Microsoft's ISA Server Web site and ISAserver.org. Tom and Deb's book on the first release of the product "Configuring ISA Server 2000" dominated the ISA Server 2000 book market having sold over 40,000 copies worldwide, and the ISA Server community is eagerly awaiting Tom and Deb's book on ISA Server 2004, which is the dramatically upgraded new release from Microsoft. This book will be featured prominently on the ISAserver.org home page as well as referenced on Microsoft TechNet and ISA Server Web pages. Tom and Deb's unparalleled technical expertise combined with prime on-line marketing opportunities will make this the #1 book again in the ISA Server market. * This book will provide readers with unparalleled information on installing, configuring, and troubleshooting ISA Server 2004 by

teaching readers to: * Deploy ISA Server 2004 in small businesses and large organizations. * Learn how to configure complex DMZ configurations using ISA Server 2004's new network awareness features and built-in multinetworking capabilities. * Learn how to take advantage of ISA Server 2004's new VPN capabilities!

The Predictive Retailer Andrew Pearson 2017-10-23

The Predictive Retailer is a retail company that utilizes the latest technological developments to deliver an exceptional personalized experience to each and every customer. Today, technology such as AI, Machine Learning, Augmented Reality, IoT, Real-time stream processing, social media, and wearables are altering the Customer Experience (CX) landscape and retailers need to jump aboard this fast moving technology or run the risk of being left out in the cold. The Predictive Retailer reveals how these and other technologies can help shape the customer journey. The book details how the five types of analytics-descriptive, diagnostic, predictive, prescriptive, and edge analytics-affect not only the customer journey, but also just about every operating function of the retailer. An IoT connected retailer can make its operations smart. Connected devices can help with inventory optimization, supply chain management, labor management, waste management, as well as keep the retailer's data centers green and its energy use smart. Social media is no longer a vanity platform,

but rather it is a place to both connect with current customers as well as court new ones. It is also a powerful branding channel that can be utilized to both understand a retailer's position in the market, as well as a place to benchmark its position against its competitors. Today, technology moves at break-neck speed and it can offer the potential of anticipatory capabilities, but it also comes with a confusing variety of technological terms--Big Data, Cognitive Computing, CX, Data Lakes, Hadoop, Kafka, Personalization, Spark, etc., etc. The Predictive Retailer will help make sense of it all, so that a retail executive can cut through the confusing technological jargon and understand why a Spark-based real-time stream processing data stream might be preferable to a TIBCO Streambase one, or an IBM Streaming Analytics one. This book will help retail executives break through the technological clutter so that they can deliver an unrivaled customer experience to each and every patron that comes through their doors.

FreeBSD Handbook FreeBSD Documentation Project
2000 The FreeBSD Handbook is a comprehensive FreeBSD tutorial and reference. It covers installation, day-to-day use of FreeBSD, and much more, such as the Ports collection, creating a custom kernel, security topics, the X Window System, how to use FreeBSD's Linux binary compatibility, and how to upgrade your system from source using the 'make world' command,

to name a few.

Liars and Outliers Bruce Schneier 2012-01-27 In today's hyper-connected society, understanding the mechanisms of trust is crucial. Issues of trust are critical to solving problems as diverse as corporate responsibility, global warming, and the political system. In this insightful and entertaining book, Schneier weaves together ideas from across the social and biological sciences to explain how society induces trust. He shows the unique role of trust in facilitating and stabilizing human society. He discusses why and how trust has evolved, why it works the way it does, and the ways the information society is changing everything.

Puppet 5 Beginner's Guide John Arundel 2017-10-05 Puppet 5 Beginner's Guide, Third Edition is a practical guide that gets you up and running with the very latest features of Puppet 5. About This Book Develop skills to run Puppet 5 on single or multiple servers without hiccups Use Puppet to create and manage cloud resources such as Amazon EC2 instances Take full advantage of powerful new features of Puppet including loops, data types, Hiera integration, and container management Who This Book Is For Puppet 5 Beginner's Guide, Third Edition is designed for those who are new to Puppet, including system administrators and developers who are looking to manage computer server systems for configuration management. No prior programming or system

administration experience is assumed. What You Will Learn Understand the latest Puppet 5 features Install and set up Puppet and discover the latest and most advanced features Configure, build, and run containers in production using Puppet's industry-leading Docker support Deploy configuration files and templates at super-fast speeds and manage user accounts and access control Automate your IT infrastructure Use the latest features in Puppet 5 onward and its official modules Manage clouds, containers, and orchestration Get to know the best practices to make Puppet more reliable and increase its performance In Detail Puppet 5 Beginner's Guide, Third Edition gets you up and running with the very latest features of Puppet 5, including Docker containers, Hiera data, and Amazon AWS cloud orchestration. Go from beginner to confident Puppet user with a series of clear, practical examples to help you manage every aspect of your server setup. Whether you're a developer, a system administrator, or you are simply curious about Puppet, you'll learn Puppet skills that you can put into practice right away. With practical steps giving you the key concepts you need, this book teaches you how to install packages and config files, create users, set up scheduled jobs, provision cloud instances, build containers, and so much more. Every example in this book deals with something real and practical that you're likely to need in your work, and you'll see the complete Puppet code that makes it happen, along

with step-by-step instructions for what to type and what output you'll see. All the examples are available in a GitHub repo for you to download and adapt for your own server setup. Style and approach This tutorial is packed with quick step-by-step instructions that are immediately applicable for beginners. This is an easy-to-read guide, to learn Puppet from scratch, that explains simply and clearly all you need to know to use this essential IT power tool, while applying these solutions to real-world scenarios.

Getting Started with Impala John Russell 2014-09-25

Learn how to write, tune, and port SQL queries and other statements for a Big Data environment, using Impala—the massively parallel processing SQL query engine for Apache Hadoop. The best practices in this practical guide help you design database schemas that not only interoperate with other Hadoop components, and are convenient for administrators to manage and monitor, but also accommodate future expansion in data size and evolution of software capabilities. Written by John Russell, documentation lead for the Cloudera Impala project, this book gets you working with the most recent Impala releases quickly. Ideal for database developers and business analysts, the latest revision covers analytics functions, complex types, incremental statistics, subqueries, and submission to the Apache incubator. Getting Started with Impala includes advice from Cloudera's development team, as well as insights from its consulting engagements with

customers. Learn how Impala integrates with a wide range of Hadoop components Attain high performance and scalability for huge data sets on production clusters Explore common developer tasks, such as porting code to Impala and optimizing performance Use tutorials for working with billion-row tables, date- and time-based values, and other techniques Learn how to transition from rigid schemas to a flexible model that evolves as needs change Take a deep dive into joins and the roles of statistics

Information Storage and Management EMC Education Services 2012-04-30 The new edition of a bestseller, now revised and update throughout! This new edition of the unparalleled bestseller serves as a full training course all in one and as the world's largest data storage company, EMC is the ideal author for such a critical resource. They cover the components of a storage system and the different storage system models while also offering essential new material that explores the advances in existing technologies and the emergence of the "Cloud" as well as updates and vital information on new technologies. Features a separate section on emerging area of cloud computing Covers new technologies such as: data de-duplication, unified storage, continuous data protection technology, virtual provisioning, FCoE, flash drives, storage tiering, big data, and more Details storage models such as Network Attached Storage (NAS), Storage Area Network (SAN), Object Based Storage along with

virtualization at various infrastructure components
Explores Business Continuity and Security in physical and virtualized environment Includes an enhanced Appendix for additional information This authoritative guide is essential for getting up to speed on the newest advances in information storage and management.

SQL Injection Attacks and Defense Justin Clarke 2012-06-18 What is SQL injection? -- Testing for SQL injection -- Reviewing code for SQL injection -- Exploiting SQL injection -- Blind SQL injection exploitation -- Exploiting the operating system -- Advanced topics -- Code-level defenses -- Platform level defenses -- Confirming and recovering from SQL injection attacks -- References.

The Night Crossing Karen Ackerman 1995-04-25 For fans of The Boy in the Striped Pajamas comes a heartstopping tale of resilience for one young girl during World War II. It's hard to leave your home and friends, but the Nazis have invaded Clara's native Austria, and her family is no longer safe. Clara and her family take only what they can carry and travel by night to the Swiss border, where they hope to escape to freedom, even while soldiers are everywhere. Can Clara's heroism carry her family across the border, with their lives and few precious possessions intact?

SAS 9.4 Intelligence Platform: Overview, Second Edition Sas Institute 2016-02-12 Provides a point of entry for understanding the basics of the SAS

Intelligence Platform. It discusses the benefits of the SAS Intelligence Platform to businesses, describes the architecture, and provides an overview of each software component. This document can be used alone or as an introduction to the deployment and administration guides.

Big Data Concepts, Theories, and Applications Shui Yu 2016-03-03 This book covers three major parts of Big Data: concepts, theories and applications. Written by world-renowned leaders in Big Data, this book explores the problems, possible solutions and directions for Big Data in research and practice. It also focuses on high level concepts such as definitions of Big Data from different angles; surveys in research and applications; and existing tools, mechanisms, and systems in practice. Each chapter is independent from the other chapters, allowing users to read any chapter directly. After examining the practical side of Big Data, this book presents theoretical perspectives. The theoretical research ranges from Big Data representation, modeling and topology to distribution and dimension reducing. Chapters also investigate the many disciplines that involve Big Data, such as statistics, data mining, machine learning, networking, algorithms, security and differential geometry. The last section of this book introduces Big Data applications from different communities, such as business, engineering and science. Big Data Concepts, Theories and Applications is designed as a reference for

researchers and advanced level students in computer science, electrical engineering and mathematics.

Practitioners who focus on information systems, big data, data mining, business analysis and other related fields will also find this material valuable.

MySQL Administrator's Guide MySQL AB 2004-07-26

If you are not already one of the 4 million plus users of the MySQL database, you will want to become one after reading this book. MySQL Administrator's Guide is the official guide to installing, administering and working with MySQL 5.0 and MySQL databases.

Starting with the basics, you will learn to store, manage and retrieve data in a MySQL database and to eventually manipulate the code behind MySQL 5.0 to customize your database. With the clear examples and tips provided in this comprehensive learning guide, you can become a MySQL expert in no time.

Big Data Bootcamp David Feinleib 2014-09-26

Investors and technology gurus have called big data one of the most important trends to come along in decades. Big Data Bootcamp explains what big data is and how you can use it in your company to become one of tomorrow's market leaders. Along the way, it explains the very latest technologies, companies, and advancements. Big data holds the keys to delivering better customer service, offering more attractive products, and unlocking innovation. That's why, to remain competitive, every organization should become a big data company. It's also why every manager and

technology professional should become knowledgeable about big data and how it is transforming not just their own industries but the global economy. And that knowledge is just what this book delivers. It explains components of big data like Hadoop and NoSQL databases; how big data is compiled, queried, and analyzed; how to create a big data application; and the business sectors ripe for big data-inspired products and services like retail, healthcare, finance, and education. Best of all, your guide is David Feinleib, renowned entrepreneur, venture capitalist, and author of *Why Startups Fail*. Feinleib's *Big Data Landscape*, a market map featured and explained in the book, is an industry benchmark that has been viewed more than 150,000 times and is used as a reference by VMWare, Dell, Intel, the U.S. Government Accountability Office, and many other organizations. Feinleib also explains:

- Why every businessperson needs to understand the fundamentals of big data or get run over by those who do
- How big data differs from traditional database management systems
- How to create and run a big data project
- The technical details powering the big data revolution

Whether you're a Fortune 500 executive or the proprietor of a restaurant or web design studio, *Big Data Bootcamp* will explain how you can take full advantage of new technologies to transform your company and your career.

Absolute FreeBSD, 2nd Edition Michael W. Lucas

2013-04-12 FreeBSD—the powerful, flexible, and free Unix-like operating system—is the preferred server for many enterprises. But it can be even trickier to use than either Unix or Linux, and harder still to master. Absolute FreeBSD, 2nd Edition is your complete guide to FreeBSD, written by FreeBSD committer Michael W. Lucas. Lucas considers this completely revised and rewritten second edition of his landmark work to be his best work ever; a true product of his love for FreeBSD and the support of the FreeBSD community. Absolute FreeBSD, 2nd Edition covers installation, networking, security, network services, system performance, kernel tweaking, filesystems, SMP, upgrading, crash debugging, and much more, including coverage of how to:

- Use advanced security features like packet filtering, virtual machines, and host-based intrusion detection
- Build custom live FreeBSD CDs and bootable flash
- Manage network services and filesystems
- Use DNS and set up email, IMAP, web, and FTP services for both servers and clients
- Monitor your system with performance-testing and troubleshooting tools
- Run diskless systems
- Manage schedulers, remap shared libraries, and optimize your system for your hardware and your workload
- Build custom network appliances with embedded FreeBSD
- Implement redundant disks, even without special hardware
- Integrate FreeBSD-specific SNMP into your network management system.

Whether you're just getting started with FreeBSD or you've been using it

for years, you'll find this book to be the definitive guide to FreeBSD that you've been waiting for.

Data Curation Cynthia Hudson-Vitale 2017 Explores the infrastructure that ARL member institutions are using for data curation, which data curation services are offered, who may use them, which disciplines demand services most, library staffing levels, policies and workflows, and the challenges of supporting these activities. Includes examples of data repository web pages, descriptions of services, infrastructure, workflows, metadata schemas, and policies, and job descriptions.

IBM PowerVM Virtualization Introduction and Configuration Scott Vetter 2015-11-24 This IBM® Redbooks® publication provides an introduction to PowerVMTM virtualization technologies on Power System servers. PowerVM is a combination of hardware, firmware, and software that provides CPU, network, and disk virtualization. These are the main virtualization technologies: POWER7, POWER6, and POWER5 hardware POWER Hypervisor Virtual I/O Server Though the PowerVM brand includes partitioning, management software, and other offerings, this publication focuses on the virtualization technologies that are part of the PowerVM Standard and Enterprise Editions. This publication is also designed to be an introduction guide for system administrators, providing instructions for these tasks: Configuration and creation of partitions and resources

on the HMC Installation and configuration of the Virtual I/O Server Creation and installation of virtualized partitions Examples using AIX, IBM i, and Linux This edition has been updated with the latest updates available and an improved content organization.

Hadoop Operations Eric Sammer 2012-09-26 If you've been asked to maintain large and complex Hadoop clusters, this book is a must. Demand for operations-specific material has skyrocketed now that Hadoop is becoming the de facto standard for truly large-scale data processing in the data center. Eric Sammer, Principal Solution Architect at Cloudera, shows you the particulars of running Hadoop in production, from planning, installing, and configuring the system to providing ongoing maintenance. Rather than run through all possible scenarios, this pragmatic operations guide calls out what works, as demonstrated in critical deployments. Get a high-level overview of HDFS and MapReduce: why they exist and how they work Plan a Hadoop deployment, from hardware and OS selection to network requirements Learn setup and configuration details with a list of critical properties Manage resources by sharing a cluster across multiple groups Get a runbook of the most common cluster maintenance tasks Monitor Hadoop clusters—and learn troubleshooting with the help of real-world war stories Use basic tools and techniques to handle backup and catastrophic failure Absolute OpenBSD, 2nd Edition Michael W. Lucas

2013-04-15 OpenBSD, the elegant, highly secure Unix-like operating system, is widely used as the basis for critical DNS servers, routers, firewalls, and more. This long-awaited second edition of Absolute OpenBSD maintains author Michael Lucas's trademark straightforward and practical approach that readers have enjoyed for years. You'll learn the intricacies of the platform, the technical details behind certain design decisions, and best practices, with bits of humor sprinkled throughout. This edition has been completely updated for OpenBSD 5.3, including new coverage of OpenBSD's boot system, security features like W^X and ProPolice, and advanced networking techniques. You'll learn how to:

- Manage network traffic with VLANs, trunks, IPv6, and the PF packet filter
- Make software management quick and effective using the ports and packages system
- Give users only the access they need with groups, sudo, and chroots
- Configure OpenBSD's secure implementations of SNMP, DHCP, NTP, hardware sensors, and more
- Customize the installation and upgrade processes for your network and hardware, or build a custom OpenBSD release

Whether you're a new user looking for a complete introduction to OpenBSD or an experienced sysadmin looking for a refresher, Absolute OpenBSD, 2nd Edition will give you everything you need to master the intricacies of the world's most secure operating system.

How to Cheat at Configuring Exchange Server 2007

Henrik Walther 2011-04-18 The Perfect Reference for the Multitasked System Administrators The new version of Exchange is an ambitious overhaul that tries to balance the growing needs for performance, cost effectiveness, and security. For the average system administrator, it will present a difficult migration path from earlier versions and a vexing number of new features. How to Cheat will help you get Exchange Server 2007 up and running as quickly and safely as possible.

- Understand Exchange 2007 Prerequisites Review hardware and software requirements, Active Directory requirements, and more.
- Manage Recipient Filtering Create a custom MMC that contains the Exchange 2007 Recipients work center, which can be used, for example, by the helpdesk staff in your organization.
- Manage Outlook Anywhere Outlook Anywhere makes it possible for your end users to remotely access their mailboxes from the Internet using their full Outlook clients.
- Manage Outlook Web Access 2007 See how Outlook Web Access 2007 was completely rewritten in managed code to make it scale even better.
- Use the Exchange 2007 Queue Viewer You can now view information about queues and examine the messages held within them.
- Master Powerful Out-of-the-Box Spam Protection The Edge Transport server supports SMTP, provides several antispam filtering agents, and supports antivirus extensibility.
- Manage a Single-Copy Cluster-Based Setup SCC-based cluster provides service failover and

still has a single point of failure when it comes to the databases. • Recover Mailbox Databases Use the improved database portability feature to port and recover a Mailbox database to any server in the Exchange 2007 organization. Essential information for the multi-tasked System Administrator charged perform everyday tasks Includes the latest coverage of the major new release of Exchange Server Emphasizes best-practice security measures for all areas and in particular the area of remote access via Outlook

Making Data Smarter with IBM Spectrum Discover: Practical AI Solutions Ivaylo B. Bozhinov 2020-10-19

More than 80% of all data that is collected by organizations is not in a standard relational database. Instead, it is trapped in unstructured documents, social media posts, machine logs, and so on. Many organizations face significant challenges to manage this deluge of unstructured data, such as the following examples: Pinpointing and activating relevant data for large-scale analytics Lacking the fine-grained visibility that is needed to map data to business priorities Removing redundant, obsolete, and trivial (ROT) data Identifying and classifying sensitive data IBM® Spectrum Discover is a modern metadata management software that provides data insight for petabyte-scale file and Object Storage, storage on-premises, and in the cloud. This software enables organizations to make better business decisions and

gain and maintain a competitive advantage. IBM Spectrum® Discover provides a rich metadata layer that enables storage administrators, data stewards, and data scientists to efficiently manage, classify, and gain insights from massive amounts of unstructured data. It improves storage economics, helps mitigate risk, and accelerates large-scale analytics to create competitive advantage and speed critical research. This IBM Redbooks® publication presents several use cases that are focused on artificial intelligence (AI) solutions with IBM Spectrum Discover. This book helps storage administrators and technical specialists plan and implement AI solutions by using IBM Spectrum Discover and several other IBM Storage products.

IBM Spectrum Discover: Metadata Management for Deep Insight of Unstructured Storage Joseph Dain 2019-10-01 This IBM® Redpaper publication provides a comprehensive overview of the IBM Spectrum® Discover metadata management software platform. We give a detailed explanation of how the product creates, collects, and analyzes metadata. Several in-depth use cases are used that show examples of analytics, governance, and optimization. We also provide step-by-step information to install and set up the IBM Spectrum Discover trial environment. More than 80% of all data that is collected by organizations is not in a standard relational database. Instead, it is trapped in unstructured documents, social media posts, machine logs, and so on. Many organizations

face significant challenges to manage this deluge of unstructured data such as: Pinpointing and activating relevant data for large-scale analytics Lacking the fine-grained visibility that is needed to map data to business priorities Removing redundant, obsolete, and trivial (ROT) data Identifying and classifying sensitive data IBM Spectrum Discover is a modern metadata management software that provides data insight for petabyte-scale file and Object Storage, storage on premises, and in the cloud. This software enables organizations to make better business decisions and gain and maintain a competitive advantage. IBM Spectrum Discover provides a rich metadata layer that enables storage administrators, data stewards, and data scientists to efficiently manage, classify, and gain insights from massive amounts of unstructured data. It improves storage economics, helps mitigate risk, and accelerates large-scale analytics to create competitive advantage and speed critical research.

Big Data Application Architecture Q&A Nitin Sawant 2014-01-24 Big Data Application Architecture Pattern Recipes provides an insight into heterogeneous infrastructures, databases, and visualization and analytics tools used for realizing the architectures of big data solutions. Its problem-solution approach helps in selecting the right architecture to solve the problem at hand. In the process of reading through these problems, you will learn harness the power of new big data opportunities which various enterprises use to

attain real-time profits. Big Data Application Architecture Pattern Recipes answers one of the most critical questions of this time 'how do you select the best end-to-end architecture to solve your big data problem?'. The book deals with various mission critical problems encountered by solution architects, consultants, and software architects while dealing with the myriad options available for implementing a typical solution, trying to extract insight from huge volumes of data in real-time and across multiple relational and non-relational data types for clients from industries like retail, telecommunication, banking, and insurance. The patterns in this book provide the strong architectural foundation required to launch your next big data application. The architectures for realizing these opportunities are based on relatively less expensive and heterogeneous infrastructures compared to the traditional monolithic and hugely expensive options that exist currently. This book describes and evaluates the benefits of heterogeneity which brings with it multiple options of solving the same problem, evaluation of trade-offs and validation of 'fitness-for-purpose' of the solution.

Data Protection Preston de Guise 2017-03-03 This is the fundamental truth about data protection: backup is dead. Or rather, backup and recovery, as a standalone topic, no longer has relevance in IT. As a standalone topic, it's been killed off by seemingly exponential growth in storage and data, by the cloud, and by

virtualization. So what is data protection? This book takes a holistic, business-based approach to data protection. It explains how data protection is a mix of proactive and reactive planning, technology and activities that allow for data continuity. It shows how truly effective data protection comes from a holistic approach considering the entire data lifecycle and all required SLAs. Data protection is neither RAID nor is it continuous availability, replication, snapshots or backups—it is all of them, combined in a considered and measured approach to suit the criticality of the data and meet all the requirements of the business. The book also discusses how businesses seeking to creatively leverage their IT investments and to drive through cost optimization are increasingly looking at data protection as a mechanism to achieve those goals. In addition to being a type of insurance policy, data protection is becoming an enabler for new processes around data movement and data processing. This book arms readers with information critical for making decisions on how data can be protected against loss in the cloud, on-premises, or in a mix of the two. It explains the changing face of recovery in a highly virtualized data center and techniques for dealing with big data. Moreover, it presents a model for where data recovery processes can be integrated with IT governance and management in order to achieve the right focus on

recoverability across the business.

High-Performance Computing in Finance M. A. H.

Dempster 2018-02-21 High-Performance Computing

(HPC) delivers higher computational performance to solve problems in science, engineering and finance.

There are various HPC resources available for different needs, ranging from cloud computing– that can be used without much expertise and expense – to more tailored hardware, such as Field-Programmable Gate Arrays (FPGAs) or D-Wave’s quantum computer systems. High-Performance Computing in Finance is the first book that provides a state-of-the-art introduction to HPC for finance, capturing both academically and practically relevant problems.

Symphony 5 Fabien Potencier 2019-11-08

Architecting Modern Data Platforms Jan Kunigk 2018-

12-05 There’s a lot of information about big data

technologies, but splicing these technologies into an end-to-end enterprise data platform is a daunting task

not widely covered. With this practical book, you’ll

learn how to build big data infrastructure both on-

premises and in the cloud and successfully architect a modern data platform. Ideal for enterprise architects,

IT managers, application architects, and data

engineers, this book shows you how to overcome the many challenges that emerge during Hadoop projects.

You’ll explore the vast landscape of tools available in

the Hadoop and big data realm in a thorough technical primer before diving into: Infrastructure: Look at all

component layers in a modern data platform, from the server to the data center, to establish a solid foundation for data in your enterprise Platform: Understand aspects of deployment, operation, security, high availability, and disaster recovery, along with everything you need to know to integrate your platform with the rest of your enterprise IT Taking Hadoop to the cloud: Learn the important architectural aspects of running a big data platform in the cloud while maintaining enterprise security and high availability

Curating Research Data Lisa R. Johnston 2016-11-01
Data are becoming the proverbial coin of the digital realm: a research commodity that might purchase reputation credit in a disciplinary culture of data sharing, or buy transparency when faced with funding agency mandates or publisher scrutiny. Unlike most monetary systems, however, digital data can flow in all too great an abundance. Not only does this currency actually grow on trees, but it comes from animals, books, thoughts, and each of us! And that is what makes data curation so essential. The abundance of digital research data challenges library and information science professionals to harness this flow of information streaming from research discovery and scholarly pursuit and preserve the unique evidence for future use. Volume One of Curating Research Data explores the variety of reasons, motivations, and drivers for why data curation services are needed in

the context of academic and disciplinary data repository efforts. Twelve chapters, divided into three parts, take an in-depth look at the complex practice of data curation as it emerges around us. Part I sets the stage for data curation by describing current policies, data sharing cultures, and collaborative efforts currently underway that impact potential services. Part II brings several key issues, such as cost recovery and marketing strategy, into focus for practitioners when considering how to put data curation services in action. Finally, Part III describes the full lifecycle of data by examining the ethical and practical reuse issues that data curation practitioners must consider as we strive to prepare data for the future. Digital data is ubiquitous and rapidly reshaping how scholarship progresses now and into the future. The information expertise of librarians can help ensure the resiliency of digital data, and the information it represents, by addressing how the meaning, integrity, and provenance of digital data generated by researchers today will be captured and conveyed to future researchers.

The Best Damn Exchange, SQL and IIS Book Period
Henrik Walther 2011-08-31 Exchange, SQL and IIS are at the core of most Microsoft enterprise servers. The 2007 releases of these products, along with the release of Windows Vista and Windows 2008 Server, represents the biggest overhaul of Windows enterprise products since Windows 2000. The dramatic changes to security tools and the addition of features that

support "anywhere access" present IT professionals with a steep learning curve. Making certain that these products are configured to meet regulatory compliance requirements adds additional complexity to day-to-day management network management. The Best Damn Exchange, SQL and IIS Book Period delivers an all-in-one reference for Windows System Administrators deploying the 2007 releases of these core Microsoft servers. The coverage is comprehensive, and provides users with just about everything they need to manage a Windows enterprise server. Special Bonus coverage includes how to gather and analyze the many log files generated by these servers. All-in-one coverage includes Exchange, SQL, and IIS Servers Integrated coverage on all key security features Bonus coverage includes analyzing server logs and integrating Communicator 2007

Implementing VxRail HCI Solutions Victor Wu 2021-06-04 Plan, design, deploy, and administer the solutions available in VxRail Appliance Key Features Learn how to plan and design the VxRail HCI system Understand VxRail's administration, lifecycle management, and cluster scale-out Explore migration methodologies for VxRail systems Book Description Hyper-converged infrastructure (HCI) can help you simplify the provisioning and daily operations of computing and storage. With this book, you'll understand how HCI can offload the day 0 deployment and day-to-day operations of a system administrator. You'll explore the

VxRail Appliance, which is an HCI solution that provides lifecycle management, automation, and operational simplicity. Starting with an overview of the VxRail Appliance system architecture and components, you'll understand the benefits of the VxRail system and compare it with the environment of traditional servers and storage. As you advance, the book covers topics such as disaster recovery and active-active and active-passive solutions for VxRail. By the end of this book, you'll have gained the confidence to manage the deployment, administration, planning, and design of a VxRail system. What you will learn

- Set up the hardware and software requirements for a VxRail installation
- Monitor the status of VxRail appliances with the VxRail Manager plugin
- Get to grips with all the administration interfaces used to manage the VxRail appliance
- Understand vCenter roles and permissions management in the VxRail cluster
- Discover best practices for vSAN configuration in the VxRail cluster
- Find out about VxRail cluster scale-out rules and how to expand the VxRail cluster
- Deploy active-passive solutions for VxRail with VMware Site Recovery Manager (SRM)

Who this book is for If you are a system architect, system administrator, or consultant involved in planning and deploying VxRail HCI or want to learn how to use VxRail HCI, then this book is for you. Equivalent knowledge and administration experience with ESXi and vCenter

Server will be helpful.

Data Analytics for Intelligent Transportation Systems

Mashrur Chowdhury 2017-04-05 Data Analytics for Intelligent Transportation Systems provides in-depth coverage of data-enabled methods for analyzing intelligent transportation systems that includes detailed coverage of the tools needed to implement these methods using big data analytics and other computing techniques. The book examines the major characteristics of connected transportation systems, along with the fundamental concepts of how to analyze the data they produce. It explores collecting, archiving, processing, and distributing the data, designing data infrastructures, data management and delivery systems, and the required hardware and software technologies. Users will learn how to design effective data visualizations, tactics on the planning process, and how to evaluate alternative data analytics for different connected transportation applications, along with key safety and environmental applications for both commercial and passenger vehicles, data privacy and security issues, and the role of social media data in traffic planning. Includes case studies in each chapter that illustrate the application of concepts covered Presents extensive coverage of existing and forthcoming intelligent transportation systems and data analytics technologies Contains contributors from both leading academic and commercial researchers Explains how to design effective data visualizations,

tactics on the planning process, and how to evaluate alternative data analytics for different connected transportation applications

Annual scientific report 1962

Using Samba Gerald Carter 2007-01-23 This book is the comprehensive guide to Samba administration, officially adopted by the Samba Team. Wondering how to integrate Samba's authentication with that of a Windows domain? How to get Samba to serve Microsoft Dfs shares? How to share files on Mac OS X? These and a dozen other issues of interest to system administrators are covered. A whole chapter is dedicated to troubleshooting! The range of this book knows few bounds. Using Samba takes you from basic installation and configuration -- on both the client and server side, for a wide range of systems -- to subtle details of security, cross-platform compatibility, and resource discovery that make the difference between whether users see the folder they expect or a cryptic error message. The current edition covers such advanced 3.x features as: Integration with Active Directory and OpenLDAP Migrating from Windows NT 4.0 domains to Samba Delegating administrative tasks to non-root users Central printer management Advanced file serving features, such as making use of Virtual File System (VFS) plugins. Samba is a cross-platform triumph: robust, flexible and fast, it turns a Unix or Linux system into a file and print server for Microsoft Windows network clients. This book will help

you make your file and print sharing as powerful and efficient as possible. The authors delve into the internals of the Windows activities and protocols to an unprecedented degree, explaining the strengths and weaknesses of each feature in Windows domains and in Samba itself. Whether you're playing on your personal computer or an enterprise network, on one note or a full three-octave range, Using Samba will give you an efficient and secure server.

Big Digital Forensic Data Darren Quick 2018-04-24

This book provides an in-depth understanding of big data challenges to digital forensic investigations, also known as big digital forensic data. It also develops the basis of using data mining in big forensic data analysis, including data reduction, knowledge management, intelligence, and data mining principles to achieve faster analysis in digital forensic investigations. By collecting and assembling a corpus of test data from a range of devices in the real world, it outlines a process of big data reduction, and evidence and intelligence extraction methods. Further, it includes the experimental results on vast volumes of real digital forensic data. The book is a valuable resource for digital forensic practitioners, researchers in big data, cyber threat hunting and intelligence, data mining and other related areas.

Microsoft Azure Security Center Yuri Diogenes 2018-06-04 Discover high-value Azure security insights, tips, and operational optimizations This book presents

comprehensive Azure Security Center techniques for safeguarding cloud and hybrid environments. Leading Microsoft security and cloud experts Yuri Diogenes and Dr. Thomas Shinder show how to apply Azure Security Center's full spectrum of features and capabilities to address protection, detection, and response in key operational scenarios. You'll learn how to secure any Azure workload, and optimize virtually all facets of modern security, from policies and identity to incident response and risk management. Whatever your role in Azure security, you'll learn how to save hours, days, or even weeks by solving problems in most efficient, reliable ways possible. Two of Microsoft's leading cloud security experts show how to:

- Assess the impact of cloud and hybrid environments on security, compliance, operations, data protection, and risk management
- Master a new security paradigm for a world without traditional perimeters
- Gain visibility and control to secure compute, network, storage, and application workloads
- Incorporate Azure Security Center into your security operations center
- Integrate Azure Security Center with Azure AD Identity Protection Center and third-party solutions
- Adapt Azure Security Center's built-in policies and definitions for your organization
- Perform security assessments and implement Azure Security Center recommendations
- Use incident response features to detect, investigate, and address threats
- Create high-fidelity fusion alerts to focus attention on

your most urgent security issues • Implement application whitelisting and just-in-time VM access • Monitor user behavior and access, and investigate compromised or misused credentials • Customize and perform operating system security baseline assessments • Leverage integrated threat intelligence to identify known bad actors

Advanced Information Systems Engineering Schahram Dustdar 2020-06-02 This book constitutes the refereed proceedings of the 32nd International Conference on Advanced Information Systems Engineering, CAiSE 2020, held in Grenoble, France, in June 2020.* The 33 full papers presented in this volume were carefully reviewed and selected from 185 submissions. The book also contains one invited talk in full paper length. The papers were organized in topical sections named: distributed applications; AI and big data in IS; process mining and analysis; requirements and modeling; and information systems engineering. Abstracts on the CAiSE 2020 tutorials can be found in the back matter of the volume. *The conference was held virtually due to the COVID-19 pandemic.

Systems Performance Brendan Gregg 2013-10-07 The Complete Guide to Optimizing Systems Performance Written by the winner of the 2013 LISA Award for Outstanding Achievement in System Administration Large-scale enterprise, cloud, and virtualized computing systems have introduced serious performance challenges. Now, internationally

renowned performance expert Brendan Gregg has brought together proven methodologies, tools, and metrics for analyzing and tuning even the most complex environments. *Systems Performance: Enterprise and the Cloud* focuses on Linux® and Unix® performance, while illuminating performance issues that are relevant to all operating systems. You'll gain deep insight into how systems work and perform, and learn methodologies for analyzing and improving system and application performance. Gregg presents examples from bare-metal systems and virtualized cloud tenants running Linux-based Ubuntu®, Fedora®, CentOS, and the illumos-based Joyent® SmartOS™ and OmniTI OmniOS®. He systematically covers modern systems performance, including the “traditional” analysis of CPUs, memory, disks, and networks, and new areas including cloud computing and dynamic tracing. This book also helps you identify and fix the “unknown unknowns” of complex performance: bottlenecks that emerge from elements and interactions you were not aware of. The text concludes with a detailed case study, showing how a real cloud customer issue was analyzed from start to finish. Coverage includes

- Modern performance analysis and tuning: terminology, concepts, models, methods, and techniques
- Dynamic tracing techniques and tools, including examples of DTrace, SystemTap, and perf
- Kernel internals: uncovering what the OS is doing
- Using system observability tools, interfaces,

and frameworks • Understanding and monitoring application performance • Optimizing CPUs: processors, cores, hardware threads, caches, interconnects, and kernel scheduling • Memory optimization: virtual memory, paging, swapping, memory architectures, busses, address spaces, and allocators • File system I/O, including caching • Storage devices/controllers, disk I/O workloads, RAID, and kernel I/O • Network-related performance issues: protocols, sockets, interfaces, and physical connections • Performance implications of OS and hardware-based virtualization, and new issues encountered with cloud computing • Benchmarking: getting accurate results and avoiding common mistakes This guide is indispensable for anyone who operates enterprise or cloud environments: system, network, database, and web admins; developers; and other professionals. For students and others new to optimization, it also provides exercises reflecting Gregg's extensive instructional experience.

OCFA Java Foundations Exam Fundamentals 1Z0-811

Hanumant Deshmukh 2020-11-11 Last Updated: 1st Jan 2021 - Build 1.1 IMPORTANT: 1. We enhance the content and publish it on Amazon regularly. We update the build number upon each release so that you can determine if a new build has been released since your purchase. If you have an older build of the eBook, please contact Amazon support to get the latest build on your device. You will then need to delete the book

from your Kindle device and then redownload it. 2.

This book DOES NOT include mock exams. It is meant to complement Enthware Mock Exams and should be used as a study guide before or while attempting the mock Exams. If you are a high schooler or a Java beginner, the 1Z0-811 exam is the best way to prove that you have learnt the basics of Java programming. OCFA Java Foundations Exam Fundamentals is a comprehensive study guide for those taking the Oracle Certified Foundations Associate, Java Certification Exam (Exam Code 1Z0-811). With complete coverage of 100% of the exam objectives, this book provides everything you need to know to take the exam confidently. Written by an expert with more than 20 years of industry experience, the book also helps you ace technical interviews by making you aware of things that technical managers focus on. The Java Foundations exam requires you to learn basic features of the language including various Java platforms, OOP, arrays, loops, decision constructs, and exception handling. This book covers all such topics thoroughly. The book also includes coding exercises that will get you moving on "write a lot of code" front. The book makes it easy to get your doubts cleared by including links to existing discussion on a particular topic. If the existing discussion doesn't address your doubt, you can see more clarification from the Author.

Storage Virtualization Tom Clark 2005 Storage virtualization has come of age, offering IT

professionals powerful new ways to simplify infrastructure, streamline management, improve utilization, and reduce costs. Now, the author of the best-selling storage books IP SANs and Designing Storage Area Networks presents an up-to-the-minute, vendor-neutral overview of storage virtualization in all its forms.

Introducing and Implementing IBM Flashsystem Ibm Redbooks 2015 IBM FlashSystem V9000 is a comprehensive all-flash enterprise storage solution that delivers the full capabilities of IBM FlashCore technology plus a rich set of software-defined storage features. It is available as one solution in a compact 6U form factor to achieve a simpler, more scalable, and cost efficient IT Infrastructure. FlashSystem V9000 improves business application availability and delivers greater resource utilization so you can get the most from your storage resources. Cases and ISV scenarios that demonstrate real-world solutions, as well as examples of the benefits gained by integrating the FlashSystem storage into business environments are discussed. --