

Access Free Emc Isilon Admin Guide Pdf File Free

The FreeBSD Handbook Dec 13 2021 "The FreeBSD Handbook" is a comprehensive FreeBSD tutorial and reference. It covers installation, day-to-day use of FreeBSD, 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.

The Predictive Airliner Jan 02 2021 The Predictive Airliner is an airline that utilizes the latest technology to deliver an exceptional personalized experience to each and every passenger it flies. Today, technology such as AI, Machine Learning, Augmented Reality, IoT, Real-time stream processing, social media, streaming analytics and wearables are altering the Customer Experience (CX) landscape and airlines need to jump aboard this fast moving technology or run the risk of being left out in the cold. The Predictive Airliner 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 operational function within an airline. An IoT-connected airline can make its operations smart. Data collected at multiple company and customer touch points can be utilized to increase customer satisfaction, as well as make the airline more profitable. The book lays out a blueprint for airlines to use to build a better overall operation. By utilizing AI, machine learning, and deep learning airlines can monitor the health of their airplanes, ensure employee satisfaction, and deliver an award-winning customer experience every time. Analytical processes like decision trees, k-means clustering, logistic regression and neural networks are explained in detail, with specific use cases detailing how they are used profitably in the aviation industry. Edge analytics, sentiment analysis, clickstream analysis, and location analysis are seen through a customer intelligence lens to ensure passengers are treated in a personalized way that will not only increase loyalty but turn passengers into apostles for the airlines they chose to fly on. Connected devices can help with inventory optimization, supply chain management, labor management, waste management, as well as keep the airline'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 an airline's position in the market, as well as a place to benchmark its position against competitors. The

Predictive Airliner reveals how airlines can utilize this channel in a multitude of ways to connect with customers, as well as help in moments of crisis. 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 Airliner will help airline executives make sense of it all, so that he or she can cut through the confusing clutter of technological jargon and understand why a Spark-based real-time stream processing data stream might be preferable to a TIBCO Streambase one, or none at all. The final chapter explains how an airline can utilize the concept of the customer journey as a roadmap to increase customer satisfaction. This book will help airline executives break through the technological clutter so that they can deliver an unrivaled customer experience to each and every passenger who steps aboard their planes.

Sams Teach Yourself Samba in 24 Hours Dec 21 2019 Author and Samba Team Member Gerald Carter dedicates most of this book to the processes of installing and configuring Samba, distinguishing between both Linux/Windows, and Unix/Windows-based systems. He discusses the mechanics of file sharing across the dual platforms and expertly outlines the means necessary to remedy common problems, including helping to identify which version of Samba you are running and where to look for the patch or bug fix you will need to resolve it. He also assists in locating utilities like autoconf that automatically compile a system configuration file with common default settings. Finally, he discusses security and password issues, as well as peripheral support for server printers.

IBM PowerVM Virtualization Introduction and Configuration Mar 24 2020
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.

MySQL Administrator's Guide Aug 09 2021 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.

Dr. Tom Shinder's Configuring ISA Server 2004 Jun 07 2021 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!

KVM Virtualization Cookbook Sep 10 2021 Deploy, manage, and scale virtual instances using Kernel-based Virtual Machines About This Book Build, manage and scale virtual machines with practical step-by-step examples Leverage the libvirt user-space tools and libraries to manage the life-cycle of KVM instances Deploy and scale applications inside KVM virtual machines with OpenStack Who This Book Is For If you are a system administrator working KVM virtualization, this book will help you grow on your expertise of working with the infrastructure to manage things in a better way. You should have a knowledge of working

with Linux based systems. What You Will Learn Deploy different workloads in isolation with KVM virtualization and better utilize the available compute resources Explore the benefits of running applications with KVM and learn to prevent the "bad-neighbor" effect Leveraging various networking technologies in the context of virtualization with Open vSwitch and the Linux bridge. Create KVM instances using Python and inspect running KVM instances Understand Kernel Tuning for enhanced KVM performance and better memory utilization In Detail Virtualization technologies such as KVM allow for better control over the available server resources, by deploying multiple virtual instances on the same physical host, or clusters of compute resources. With KVM it is possible to run various workloads in isolation with the hypervisor layer providing better tenant isolation and higher degree of security. This book will provide a deep dive into deploying KVM virtual machines using qemu and libvirt and will demonstrate practical examples on how to run, scale, monitor, migrate and backup such instances. You will also discover real production ready recipes on deploying KVM instances with OpenStack and how to programatically manage the life cycle of KVM virtual machines using Python. You will learn numerous tips and techniques which will help you deploy & plan the KVM infrastructure. Next, you will be introduced to the working of libvirt libraries and the iPython development environment. Finally, you will be able to tune your Linux kernel for high throughput and better performance. By the end of this book, you will gain all the knowledge needed to be an expert in working with the KVM virtualization infrastructure. Style and approach This book takes a complete practical approach with many step-by-step example recipes on how to use KVM in production. The book assumes certain level of expertise with Linux systems and virtualization in general. Some knowledge of Python programming is encouraged, to fully take advantage of the code recipes.

IBM PowerVC Version 2.0 Introduction and Configuration

Mar 04 2021

IBM® Power Virtualization Center (IBM® PowerVCTM) is an advanced enterprise virtualization management offering for IBM Power Systems. This IBM Redbooks® publication introduces IBM PowerVC and helps you understand its functions, planning, installation, and setup. It also shows how IBM PowerVC can integrate with systems management tools such as Ansible or Terraform and that it also integrates well into a OpenShift container environment. IBM PowerVC Version 2.0.0 supports both large and small deployments, either by managing IBM PowerVM® that is controlled by the Hardware Management Console (HMC), or by IBM PowerVM NovaLink. With this capability, IBM PowerVC can manage IBM AIX®, IBM i, and Linux workloads that run on IBM POWER® hardware. IBM PowerVC is available as a Standard Edition, or as a Private Cloud Edition. IBM PowerVC includes the following features and benefits: Virtual image capture, import, export, deployment, and management

Policy-based virtual machine (VM) placement to improve server usage
Snapshots and cloning of VMs or volumes for backup or testing purposes
Support of advanced storage capabilities such as IBM SVC vdisk
mirroring of IBM Global Mirror Management of real-time optimization
and VM resilience to increase productivity VM Mobility with placement
policies to reduce the burden on IT staff in a simple-to-install and
easy-to-use graphical user interface (GUI) Automated Simplified Remote
Restart for improved availability of VMs ifor when a host is down Role-
based security policies to ensure a secure environment for common
tasks The ability to enable an administrator to enable Dynamic
Resource Optimization on a schedule IBM PowerVC Private Cloud Edition
includes all of the IBM PowerVC Standard Edition features and
enhancements: A self-service portal that allows the provisioning of
new VMs without direct system administrator intervention. There is an
option for policy approvals for the requests that are received from
the self-service portal. Pre-built deploy templates that are set up by
the cloud administrator that simplify the deployment of VMs by the
cloud user. Cloud management policies that simplify management of
cloud deployments. Metering data that can be used for chargeback. This
publication is for experienced users of IBM PowerVM and other
virtualization solutions who want to understand and implement the next
generation of enterprise virtualization management for Power Systems.
Unless stated otherwise, the content of this publication refers to IBM
PowerVC Version 2.0.0.

Making Data Smarter with IBM Spectrum Discover: Practical AI
Solutions _ Mar 16 2022 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.

Using Samba May 06 2021 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.

Building a SAN-less Private Cloud with IBM PowerVM and IBM PowerVC Sep 29 2020 This IBM® Redpaper™ publication describes a software-defined infrastructure (SDI) solution with IBM PowerVC. In IBM PowerVC SDI, you combine scale-out IBM Power Systems™ servers with software that creates the whole stack that is needed to provide virtual machines (VMs) for applications such as open source databases or Hadoop. The SDI solution uses base IBM Power Systems technologies such as IBM PowerVM® NovaLink and the open source hypervisor kernel-based virtual machine (KVM). The solution combines with sophisticated storage technologies, such as IBM Spectrum™ Scale, and with the powerful networking capabilities that are provided by the Open vSwitch (OVS) technology. IBM PowerVC "hides" much of this software so that it is not apparent to your daily cloud operations. By using IBM PowerVC, you can manage scale-out SDI-based systems along with traditional PowerVM systems. This publication describes how to install and configure the SDI solution using PowerVM and IBM PowerVC running on the Power Systems platform. This publication also presents the

essentials to help existing Power Systems technical specialists use existing "under the covers" disk space to build a cost-effective cloud solution.

Big Digital Forensic Data Jan 14 2022 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.

Data Curation Jan 22 2020 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.

Data Protection Feb 27 2023 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.

Secondary Social Studies for Pakistan _____ Dec 01 2020

Architecting Modern Data Platforms _____ Mar 28 2023 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

LDAP System Administration _____ Jun 26 2020 Be more productive and make your life easier. That's what LDAP System Administration is all about. System administrators often spend a great deal of time managing configuration information located on many different machines: usernames, passwords, printer configurations, email client configurations, and network filesystem configurations, to name a few. LDAPv3 provides tools for centralizing all of the configuration information and placing it under your control. Rather than maintaining several administrative databases (NIS, Active Directory, Samba, and NFS configuration files), you can make changes in only one place and have all your systems immediately "see" the updated information. Practically platform independent, this book uses the widely available, open source OpenLDAP 2 directory server as a premise for examples, showing you how to use it to help you manage your configuration information effectively and securely. OpenLDAP 2 ships with most Linux® distributions and Mac OS® X, and can be easily downloaded for most Unix-based systems. After introducing the workings of a directory service and the LDAP protocol, all aspects of building and installing OpenLDAP, plus key ancillary packages like SASL and OpenSSL, this book discusses: Configuration and access control Distributed directories; replication and referral Using OpenLDAP to replace NIS Using OpenLDAP to manage email configurations Using LDAP

for abstraction with FTP and HTTP servers, Samba, and Radius Interoperating with different LDAP servers, including Active Directory Programming using Net::LDAP If you want to be a master of your domain, LDAP System Administration will help you get up and running quickly regardless of which LDAP version you use. After reading this book, even with no previous LDAP experience, you'll be able to integrate a directory server into essential network services such as mail, DNS, HTTP, and SMB/CIFS.

Rosie Revere, Engineer Feb 03 2021 New York Times Bestseller Rosie may seem quiet during the day, but at night she's a brilliant inventor of gizmos and gadgets who dreams of becoming a great engineer. When her great-great-aunt Rose (Rosie the Riveter) comes for a visit and mentions her one unfinished goal—to fly—Rosie sets to work building a contraption to make her aunt's dream come true. But when her contraption doesn't fly but rather hovers for a moment and then crashes, Rosie deems the invention a failure. On the contrary, Aunt Rose insists that Rosie's contraption was a raging success: you can only truly fail, she explains, if you quit. From the powerhouse author-illustrator team of Iggy Peck, Architect comes Rosie Revere, Engineer, another charming, witty picture book about believing in yourself and pursuing your passion. Ada Twist, Scientist, the companion picture book featuring the next kid from Iggy Peck's class, is available in September 2016.!--?xml:namespace prefix = o ns = "urn:schemas-microsoft-com:office:office" /-- Praise for Rosie Revere, Engineer "Comically detailed mixed-media illustrations that keep the mood light and emphasize Rosie's creativity at every turn."—Publishers Weekly "The detritus of Rosie's collections is fascinating, from broken dolls and stuffed animals to nails, tools, pencils, old lamps and possibly an erector set. And cheddar-cheese spray." —Kirkus Reviews "This celebration of creativity and perseverance is told through rhyming text, which gives momentum and steady pacing to a story, consistent with the celebration of its heroine, Rosie. She's an imaginative thinker who hides her light under a bushel (well, really, the bed) after being laughed at for one of her inventions." —Booklist Award 2013 Parents' Choice Award - GOLD 2014 Amelia Bloomer Project List ReadBoston's Best Read Aloud Book

High-Performance Computing in Finance Nov 12 2021 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.

Microsoft Azure Security Center Aug 29 2020 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

Plunkett's Companion to the Almanac of American Employers 2009 2023 Contains profiles of hundreds of the best, rapidly-growing mid-size employers of 100 to 2,500 employees. These are highly-successful companies, located nationwide, that are of vital importance to job-seekers of all types.

Jan 26

Annual Scientific Report Feb 21 2020

Puppet 5 Beginner's Guide Apr 24 2020 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.

[The Predictive Retailer](#) Jul 08 2021 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.

Curating Research Data Oct 11 2021 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.

Hadoop Operations Jun 19 2022 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

Data Analytics for Intelligent Transportation Systems

Apr 17 2022

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

Implementing VxRail HCI Solutions Sep 22 2022 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.

Introduction to Storage Area Networks Apr 05 2021

The superabundance of data that is created by today's businesses is making storage a strategic investment priority for companies of all sizes. As storage takes precedence, the following major initiatives emerge:

Flatten and converge your network: IBM® takes an open, standards-based approach to implement the latest advances in the flat, converged data center network designs of today. IBM Storage solutions enable clients to deploy a high-speed, low-latency Unified Fabric Architecture.

Optimize and automate virtualization: Advanced virtualization awareness reduces the cost and complexity of deploying physical and virtual data center infrastructure.

Simplify management: IBM data center networks are easy to deploy, maintain, scale, and virtualize, delivering the foundation of consolidated operations for dynamic infrastructure management.

Storage is no longer an afterthought. Too much is at stake. Companies are searching for more ways to efficiently manage expanding volumes of data, and to make that data accessible throughout the enterprise. This demand is propelling the move of storage into the network. Also, the increasing complexity of managing large numbers of storage devices and

vast amounts of data is driving greater business value into software and services. With current estimates of the amount of data to be managed and made available increasing at 60% each year, this outlook is where a storage area network (SAN) enters the arena. SANs are the leading storage infrastructure for the global economy of today. SANs offer simplified storage management, scalability, flexibility, and availability; and improved data access, movement, and backup. Welcome to the cognitive era. The smarter data center with the improved economics of IT can be achieved by connecting servers and storage with a high-speed and intelligent network fabric. A smarter data center that hosts IBM Storage solutions can provide an environment that is smarter, faster, greener, open, and easy to manage. This IBM® Redbooks® publication provides an introduction to SAN and Ethernet networking, and how these networks help to achieve a smarter data center. This book is intended for people who are not very familiar with IT, or who are just starting out in the IT world.

Hadoop Application Architectures Oct 23 2022 Get expert guidance on architecting end-to-end data management solutions with Apache Hadoop. While many sources explain how to use various components in the Hadoop ecosystem, this practical book takes you through architectural considerations necessary to tie those components together into a complete tailored application, based on your particular use case. To reinforce those lessons, the book's second section provides detailed examples of architectures used in some of the most commonly found Hadoop applications. Whether you're designing a new Hadoop application, or planning to integrate Hadoop into your existing data infrastructure, Hadoop Application Architectures will skillfully guide you through the process. This book covers: Factors to consider when using Hadoop to store and model data Best practices for moving data in and out of the system Data processing frameworks, including MapReduce, Spark, and Hive Common Hadoop processing patterns, such as removing duplicate records and using windowing analytics Giraph, GraphX, and other tools for large graph processing on Hadoop Using workflow orchestration and scheduling tools such as Apache Oozie Near-real-time stream processing with Apache Storm, Apache Spark Streaming, and Apache Flume Architecture examples for clickstream analysis, fraud detection, and data warehousing

Managing the Mail Jul 20 2022

Information Storage and Management Dec 25 2022 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.

Virtualizing Hadoop Apr 29 2023 Plan and Implement Hadoop Virtualization for Maximum Performance, Scalability, and Business Agility Enterprises running Hadoop must absorb rapid changes in big data ecosystems, frameworks, products, and workloads. Virtualized approaches can offer important advantages in speed, flexibility, and elasticity. Now, a world-class team of enterprise virtualization and big data experts guide you through the choices, considerations, and tradeoffs surrounding Hadoop virtualization. The authors help you decide whether to virtualize Hadoop, deploy Hadoop in the cloud, or integrate conventional and virtualized approaches in a blended solution. First, *Virtualizing Hadoop* reviews big data and Hadoop from the standpoint of the virtualization specialist. The authors demystify MapReduce, YARN, and HDFS and guide you through each stage of Hadoop data management. Next, they turn the tables, introducing big data experts to modern virtualization concepts and best practices. Finally, they bring Hadoop and virtualization together, guiding you through the decisions you'll face in planning, deploying, provisioning, and managing virtualized Hadoop. From security to multitenancy to day-to-day management, you'll find reliable answers for choosing your best Hadoop strategy and executing it. Coverage includes the following:

- Reviewing the frameworks, products, distributions, use cases, and roles associated with Hadoop
- Understanding YARN resource management, HDFS storage, and I/O
- Designing data ingestion, movement, and organization for modern enterprise data platforms
- Defining SQL engine strategies to meet strict SLAs
- Considering security, data isolation, and scheduling for multitenant environments
- Deploying Hadoop as a service in the cloud
- Reviewing the essential concepts, capabilities, and terminology of virtualization
- Applying current best practices, guidelines, and key metrics for Hadoop virtualization
- Managing multiple Hadoop frameworks and products as one unified system
- Virtualizing master and worker nodes to maximize availability and performance
- Installing and configuring Linux for a Hadoop environment

IBM Spectrum Discover: Metadata Management for Deep Insight of Unstructured Storage Aug 21 2022 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.

OCFA Java Foundations Exam Fundamentals 1Z0-811 Feb 15 2022 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.

CELPIP Practice Oct 31 2020 CELPIP practice questions, prepared by a dedicated team of exam experts, with full downloadable audio, detailed answer key, exam tips and multiple choice strategies! Here is what the CELPIP Practice Package can do for you: Practice Tests are the best way to prepare for an exam and this is the book that you need to fully prepare for the CELPIP. Here are 2 sets of CELPIP questions, prepared by a dedicated team of experts that will prepare you for the exam like nothing else will. Practice Tests familiarize you with the exam format and types of questions, giving you more confidence when you take the exam. Practice tests are a critical self-assessment tool that reveals your strengths and weaknesses. Practice tests allow you to practice your exam time management - a critical exam-writing skill that can easily improve your grade substantially. Practice tests reduce Test Anxiety, one of the main reasons for low marks on an exam. Hundreds of questions with detailed solutions and explanations to improve your understand of the basic concepts behind the questions. If you are taking the Canadian English Language Proficiency Index test - the Practice the CELPIP is an essential part of studying and passing! Each set of CELPIP Questions: Reading - Reading comprehension questions with full answer key and detailed step-by-step answer key with test tips and multiple choice strategies! Listening - Includes FULL AUDIO for listening comprehension practice - use easy scan QR codes for downloadable audio! Writing - Detailed instructions with walk-through examples and step-by-step help! Speaking - Detailed examples for all 8 speaking questions. Practice tests are a critical self-assessment tool that reveals your strengths and weaknesses familiarize you with the exam format and types of questions, build your self confidence, and practice your exam time management. All of these can make a huge difference in your score! Practice Tests also reduce Test Anxiety, one of the main reasons for low marks on an exam. Why not do everything you can to get the best score on the CELPIP?

Getting Started with Impala Nov 24 2022 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

Systems Performance Jul 28 2020 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.

The A.I. Marketer _____ May 18 2022 We seem to be living in the age of A.I. Everywhere you look, companies are touting their most recent A.I., machine learning, and deep learning breakthroughs, even when they are far short of anything that could be touted as a "breakthrough." "A.I." has eclipsed "Blockchain" and "Crypto" as the buzzword of today. Indeed, one of the best ways to raise VC funding is to stick 'AI' or 'ML' at the front of your prospectus and ".ai" at the end of your website. Separating fact from fiction is more important than it has ever been. The A.I. Marketer breaks down A.I., machine learning, and deep learning into five unique use cases—sound, time series, text, image, and video—and also reveals how marketing executives can utilize this powerful technology to help them more finely tune their marketing campaigns, better segment their customers, increase lead generation, and foster strong customer loyalty. Today, "Personalization"—the process of utilizing mobile, social, geo-location data, web morphing, context and even affective computing to tailor messages and experiences to an individual interacting with them—is becoming the optimum word in a radically new customer intelligence environment. The A.I. Marketer explains this complex technology in simple to understand terms and then shows how marketers can utilize the psychology of personalization with A.I. to both create more effective marketing campaigns as well as increase customer loyalty. Pearson shows companies how to avoid Adobe's warning of not using industrial-age technology in the digital era. Pearson also reveals how to create a platform of technology that seamlessly integrates EDW and real-time streaming data with social media content. Analytical models and neural nets can then be built on both commercial and open source technology to better understand the customer, thereby strengthening the brand and, just as importantly, increasing ROI.

How to Cheat at Configuring Exchange Server 2007 _____ May 26 2020 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

meet.uninter.edu.py