

## Isilon Administration Guide

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

Continuing its commitment to developing and delivering industry-leading storage technologies, IBM is introducing the IBM Real-time Compression Appliances for NAS, an innovative new storage offering that delivers essential storage efficiency technologies, combined with exceptional ease of use and performance. In an era when the amount of information, particularly in unstructured files, is exploding, but budgets for storing that information are stagnant, IBM Real-time Compression technology offers a powerful tool for better information management, protection, and access. IBM Real-time Compression can help slow the growth of storage acquisition, reducing storage costs while simplifying both operations and management. It also enables organizations to keep more data available for use rather than storing it offsite or on harder-to-access tape, so they can support improved analytics and decision making. IBM Real-time Compression Appliances provide on-line storage optimization through real-time data compression, delivering dramatic cost reduction without performance degradation. This IBM® Redbooks® publication is an easy-to-follow guide that describes how to design solutions successfully using IBM Real-time Compression Appliances (IBM RTCAs). It provides practical installation examples, ease of use, remote management, high availability, and administration techniques. Furthermore, it explains best practices for RTCA solution design, application integration, and practical RTCA use cases.

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.

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.

The Linux Programming Interface (TLPI) is the definitive guide to the Linux and UNIX programming interface—the interface employed by nearly every application that runs on a Linux or UNIX system. In this authoritative work, Linux programming expert Michael Kerrisk provides detailed descriptions of the system calls and library functions that you need in order to master the craft of system programming, and accompanies his explanations with clear, complete example programs. You'll find descriptions of over 500 system calls and library functions, and more than 200 example programs, 88 tables, and 115 diagrams. You'll learn how to:

- Read and write files efficiently
- Use signals, clocks, and timers
- Create processes and execute programs
- Write secure programs
- Write multithreaded programs using POSIX threads
- Build and use shared libraries
- Perform interprocess communication using pipes, message queues, shared memory, and semaphores
- Write network applications with the sockets API

While The Linux Programming Interface covers a wealth of Linux-specific features, including epoll, inotify, and the /proc file system, its emphasis on UNIX standards (POSIX.1-2001/SUSv3 and POSIX.1-2008/SUSv4) makes it equally valuable to programmers working on other UNIX platforms. The Linux Programming Interface is the most comprehensive single-volume work on the Linux and UNIX programming interface, and a book that's destined to become a new classic.

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.

"The expert's voice in big data"--Cover.

Filesystems for everything! FreeBSD includes many special-purpose filesystems. FreeBSD Mastery: Specialty Filesystems takes you through these filesystems, helping you solve problems you didn't know you have. These filesystems underlie everything from

application servers to jails. You'll learn to manage: -removable media -MSDOS and Linux filesystems -CD, DVD, and Blu-Ray disks -devfs -process filesystems -the file descriptor filesystem -the POSIX message queue filesystem -union mounts -null mounts -both memory filesystems, and when to use each -NFS versions 2, 3, and 4 -iSCSI targets and initiators -the Common Internet File System -networked disk redundancy with geom\_gate and HAST -high availability storage with HAST and iSCSI -NFSv4 Access Control Lists -User Space Filesystems -FreeBSD's Solaris-compatible automounter System administrators of any expertise level will expand their FreeBSD mastery with *FreeBSD Mastery: Specialty Filesystems*.

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.

"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® SmartOSTM 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."--Back cover.

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.

Dr Mike Moreno's 17 Day Diet is a revolutionary new weight-loss programme that activates your skinny gene so that you burn fat day in and day out. The diet is structured around four 17-day cycles: Accelerate- the rapid weight loss portion that helps flush sugar and fat storage from your system; Activate-the metabolic restart portion with alternating low and high calorie days to help shed body fat; Achieve - this phase is about learning to control portions and introducing new fitness routines; Arrive - A combination of the first three cycles to keep good habits up for good. Each cycle changes your calorie count and the food that you're eating. The variation that Dr. Mike calls 'body confusion' is designed to keep your metabolism guessing. This is not a diet that relies on a tiny list of approved foods, gruelling exercise routines, or unrealistic calorie counts that leave you hungry and unfulfilled. Each phase comes with extensive lists of what dieters can and can't eat while on the phase, but also offers acceptable cheats. He advises readers not to drink while on the diet, but concedes that if they absolutely have to then they should at least drink red wine. Dr Mike knows that a diet can only work if it's compatible with the real world, and so he's designed the programme with usability as a top priority.

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

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.

Fundamental methods and applications; Fundamental theory and further methods;

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.

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.

This IBM® Redpaper publication explains how IBM Spectrum® Discover integrates with the IBM Watson® Knowledge Catalog (WKC) component of IBM Cloud® Pak for Data (IBM CP4D) to make the enriched catalog content in IBM Spectrum Discover along with the associated data available in WKC and IBM CP4D. From an end-to-end IBM solution point of view, IBM CP4D and WKC provide state-of-the-art data governance, collaboration, and artificial intelligence (AI) and analytics tools, and IBM Spectrum Discover complements these features by adding support for unstructured data on large-scale file and object storage systems on premises and in the cloud. Many organizations face challenges to manage unstructured data. Some challenges that companies face include: Pinpointing and activating relevant data for large-scale analytics, machine learning (ML) and deep learning (DL) workloads. Lacking the fine-grained visibility that is needed to map data to business priorities. Removing redundant, obsolete, and trivial (ROT) data and identifying data that can be moved to a lower-cost storage tier. Identifying and classifying sensitive data as it relates to various compliance mandates, such as the General Data Privacy Regulation (GDPR), Payment Card Industry Data Security Standards (PCI-DSS), and the Health Information Portability and Accountability Act (HIPAA). This paper describes how IBM Spectrum Discover provides seamless integration of data in IBM Storage with IBM Watson Knowledge Catalog (WKC). Features include: Event-based cataloging and tagging of unstructured data across the enterprise. Automatically inspecting and classifying over 1000 unstructured data types, including genomics and imaging specific file formats. Automatically registering assets with WKC based on IBM Spectrum Discover search and filter criteria, and by using assets in IBM CP4D. Enforcing data governance policies in WKC in IBM CP4D based on insights from IBM Spectrum Discover, and using assets in IBM CP4D. Several in-depth use cases are used that show examples of healthcare, life sciences, and financial services. IBM Spectrum Discover integration with WKC enables storage administrators, data stewards, and data scientists to efficiently manage, classify, and gain insights from massive amounts of data. The integration improves storage economics, helps mitigate risk, and accelerates large-scale analytics to create competitive advantage and speed critical research.

This collection explores the expansion of Chinese outbound investments, aimed to sustain the increased need for natural resources, and how they have amplified the magnitude of a possible international crisis that the People's Republic of China may face in the near future by bringing together the views of a wide range of scholars. President Xi's Belt and Road initiative (BRI), aimed to promote economic development and exchanges with China for over 60 countries,

necessitates a wide range of security procedures. While the threats to Chinese enterprises and Chinese workers based on foreign soil are poised to increase, there is an urgent need to develop new guidelines for risk assessment, special insurance and crisis management. While the Chinese State Owned Enterprises are expanding their international reach capabilities, they still do not have the capacity to assure adequate security. In such a climate, this collection will be of profound value to policy makers, those working in the financial sector, and academics.

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.

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.

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. "The book that Microsoft should have written, but didn't." --Jeremy Allison, Samba Team "Your detailed explanations are clear and backed-up with source code--and the numerous bits of humor make a dry subject very enjoyable to read."

--J.D. Lindemann, network engineer, Adaptec, Inc. The first developer's guide to Microsoft(R)'s Internet/Intranet file sharing standard For years, developers and administrators have struggled to understand CIFS, Microsoft's poorly documented standard for Internet file sharing. Finally, there is an authoritative, cross-platform guide to CIFS capabilities and behavior. Implementing CIFS not only delivers the priceless knowledge of a Samba Team member dedicated to investigating the inner workings of CIFS, it also identifies and describes crucial specifications and supporting documents. Provides essential information for designing and debugging large Windows(R) and/or Samba networks Offers clear, in-depth introductions to Server Message Block (SMB), NetBIOS over TCP/IP (NBT), browser services, and authentication Drills down into the internals of CIFS, exposing its behavior on the wire and at the desktop--and its strange quirks Presents illustrative code examples throughout Reflects years of work reviewing obscure documentation, packet traces, and sourcecode Includes the SNIA CIFS Technical Reference Implementing CIFS will be indispensable to every developer who wants to provide CIFS compatibility--and every administrator or security specialist who needs an in-depth understanding of how it really works.

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

This IBM® Redbooks® publication provides a documented deployment model for IBM GPFS™ in a cross-platform environment with IBM Power Systems™, Linux, and Windows servers. With IBM GPFS, customers can have a planned foundation for file systems management for cross-platform access solutions. This book examines the functional, integration, simplification, and usability changes with GPFS v3.4. It can help the technical teams provide file system management solutions and technical support with GPFS, based on Power Systems virtualized environments for cross-platform file systems management. The book provides answers to your complex file systems management requirements, helps you maximize file system availability, and provides expert-level documentation to transfer the how-to skills to the worldwide support teams. The audience for this book is the technical professional (IT consultants, technical support staff, IT architects, and IT specialists) who is responsible for providing file system management solutions and support for cross-platform environments that are based primarily on Power Systems.

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.

This IBM® Redbooks® publication provides a practical guide to the design, installation, configuration, and maintenance of IBM Content Manager OnDemand Version 9.5. Content Manager OnDemand manages the high-volume storage and retrieval of electronic statements and provides efficient enterprise report management. Content Manager OnDemand transforms formatted computer output and printed reports, such as statements and invoices, into electronic information for easy report management. Content Manager OnDemand helps eliminate costly, high-volume print output by capturing, indexing, archiving, and presenting electronic information for improved customer service. This publication covers the key areas of Content Manager OnDemand, some of which might not be known to the Content Manager OnDemand community or are misunderstood. The book covers various topics, including basic information in administration, database structure, storage management, and security. In addition, the book covers data indexing, loading, conversion, and expiration. Other topics include user exits, performance, retention management, records management, and many more. Because many other resources are available that address subjects on different platforms, this publication is not intended as a comprehensive guide for Content Manager OnDemand. Rather, it is intended to complement the existing Content Manager OnDemand documentation and provide insight into the issues that might be encountered in the setup and use of Content Manager OnDemand. This book is intended for individuals who need to design, install, configure, and maintain Content Manager OnDemand.

Understand and implement VMware Virtual SAN: the heart of tomorrow's Software-Defined Datacenter (SDDC)  
VMware's breakthrough Software-Defined Datacenter (SDDC) initiative can help you virtualize your entire datacenter: compute, storage, networks, and associated services. Central to SDDC is VMware Virtual SAN (VSAN): a fully distributed storage architecture seamlessly integrated into the hypervisor and capable of scaling to meet any enterprise storage requirement. Now, the leaders of VMware's wildly popular Virtual SAN previews have written the first authoritative guide to this pivotal technology. You'll learn what Virtual SAN is, exactly what it offers, how to implement it, and how to maximize its value. Writing for administrators, consultants, and architects, Cormac Hogan and Duncan Epping show how Virtual SAN implements both object-based storage and a policy platform that simplifies VM storage placement. You'll learn how Virtual SAN and vSphere work together to dramatically improve resiliency, scale-out storage functionality, and control over QoS. Both an up-to-the-minute reference and hands-on tutorial, Essential Virtual SAN uses realistic examples to demonstrate Virtual SAN's most powerful capabilities. You'll learn how to plan, architect, and deploy Virtual SAN successfully, avoid gotchas, and troubleshoot problems once you're up and running. Coverage includes

Understanding the key goals and concepts of Software-Defined Storage and Virtual SAN technology Meeting physical and virtual requirements for safe Virtual SAN implementation Installing and configuring Virtual SAN for your unique environment Using Storage Policy Based Management to control availability, performance, and reliability Simplifying deployment with VM Storage Policies Discovering key Virtual SAN architectural details: caching I/O, VASA, witnesses, pass-through RAID, and more Ensuring efficient day-to-day Virtual SAN management and maintenance Interoperating with other VMware features and products Designing and sizing Virtual SAN clusters Troubleshooting, monitoring, and performance optimization

Need to master Avid Liquid 7 fast? Try a Visual QuickPro! This best-selling reference's visual format and step-by-step, task-based instructions will have you up and running with this popular editing software in no time. In this completely updated edition of our best-selling guide to Avid Liquid (formerly known as Pinnacle Liquid), leading Avid evangelist and filmmaking consultant Paul Ekert uses crystal-clear instructions and friendly prose to introduce you everything that's new in Avid Liquid 7. Filled with step-by-step, task-based instructions and loads of visual aids, this book explains how to effectively use Avid Liquid 7 to streamline your workflow; enhance your creativity; use primary functions such as Timeline, FX, Titles, Audio, DVD creation; and more. Along the way, you'll find extensive coverage of all the program's new features: SmartSound (enables you to quickly create soundtracks from a library of samples); Real-time Time Warp (see your slow-mo without rendering); 50 new special effects; dozens of style and compositing effects taken from Commotion; and more! • Takes a visual, task-based approach to guiding you through advanced topics and applying what you learn • Works like a reference book—you look up what you need and then get straight to work • Concise, straightforward steps and explanations offer the fastest way to learn tasks and concepts

Offers a Ruby tutorial featuring fifty-two exercises that cover such topics as installing the Ruby environment, organizing and writing code, strings and text, object-oriented programming, debugging and automated testing, and basic game development.

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

Although there are already some books published on Big Data, most of them only cover basic concepts and society impacts and ignore the internal implementation details-making them unsuitable to R&D people. To fill such a need, *Big Data: Storage, Sharing, and Security* examines Big Data management from an R&D perspective. It covers the 3S design. We hope you'll enjoy our Version 6.0 Blank Lined Journal in the standard size 6 x 9 inch; 15.24 x 22.86 cm as much as we did creating it for you. Here is a beautiful portable journal suitable for every 6 year old. Journal features include: 120 white pages. Gorgeous designed cover. Standard size 6 x 9 inch: 15.24 x 22.86 cm dimensions; the ideal size for all purposes fitting perfectly into your back pack or satchel. The bold white paper is sturdy enough to be used with fountain pens. Reliable standards Book industry perfect binding (the same standard binding as the books in your local library). Tough glossy paperback. Crisp white paper, with quality that minimizes ink bleed-through. The book is great for either pen or pencil users. Journals are the perfect gift for the birthdays. Click The Buy Button At The Bottom Of The Page To Begin

Journal features include: 120 white pages. Gorgeous designed cover. Standard size 6 x 9 inch: 15.24 x 22.86 cm dimensions; the ideal size for all purposes fitting perfectly into your back pack or satchel. The bold white paper is sturdy enough to be used with fountain pens. Reliable standards Book industry perfect binding (the same standard binding as the books in your local library). Tough glossy paperback. Crisp white paper, with quality that minimizes ink bleed-through. The book is great for either pen or pencil users. Journals are the perfect gift for the birthdays. Click The Buy Button At The Bottom Of The Page To Begin

Businesses of all sizes are faced with the challenge of managing huge volumes of data that are becoming increasingly valuable. But storing this data can be costly, and extracting value from the data is becoming more and more difficult. IT organizations have limited resources and cannot afford to make investment mistakes. The IBM® Storwize® V3500 system provides a smarter solution that is affordable, simple, and efficient, which enables businesses to overcome their storage challenges. IBM Storwize V3500 is the most recent addition to the IBM Storwize family of disk systems. It delivers easy-to-use, entry-level configurations that are specifically designed to meet the modest budgets

of small and medium-sized businesses. IBM Storwize V3500 features the following highlights: - Consolidate and share data with low cost iSCSI storage networking. - Deploy storage in minutes and perform storage management tasks quickly and easily through a breakthrough graphical user interface. - Experience peace of mind with proven IBM Storwize family high-availability data protection with snapshot technology and IBM warranty support. - Optimize efficiency by allocating only the amount of disk space needed at the time it is required with high performance, thin-provisioning capabilities.

“Dictionary, n: A malevolent literary device for cramping the growth of a language and making it hard and inelastic. This dictionary, however, is a most useful work.” Bierce’s groundbreaking Devil’s Dictionary had a complex publication history. Started in the mid-1800s as an irregular column in Californian newspapers under various titles, he gradually refined the new-at-the-time idea of an irreverent set of glossary-like definitions. The final name, as we see it titled in this work, did not appear until an 1881 column published in the periodical The San Francisco Illustrated Wasp. There were no publications of the complete glossary in the 1800s. Not until 1906 did a portion of Bierce’s collection get published by Doubleday, under the name The Cynic’s Word Book—the publisher not wanting to use the word “Devil” in the title, to the great disappointment of the author. The 1906 word book only went from A to L, however, and the remainder was never released under the compromised title. In 1911 the Devil’s Dictionary as we know it was published in complete form as part of Bierce’s collected works (volume 7 of 12), including the remainder of the definitions from M to Z. It has been republished a number of times, including more recent efforts where older definitions from his columns that never made it into the original book were included. Due to the complex nature of copyright, some of those found definitions have unclear public domain status and were not included. This edition of the book includes, however, a set of definitions attributed to his one-and-only “Demon’s Dictionary” column, including Bierce’s classic definition of A: “the first letter in every properly constructed alphabet.” Bierce enjoyed “quoting” his pseudonyms in his work. Most of the poetry, dramatic scenes and stories in this book attributed to others were self-authored and do not exist outside of this work. This includes the prolific Father Gassalasca Jape, whom he thanks in the preface—“jape” of course having the definition: “a practical joke.” This book is a product of its time and must be approached as such. Many of the definitions hold up well today, but some might be considered less palatable by modern readers. Regardless, the book’s humorous style is a valuable snapshot of American culture from past centuries. This book is part of the Standard Ebooks project, which produces free public domain ebooks.

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.

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.

[Copyright: 7f2200bc7396323c8a195aaa2f1fc34d](#)