
Read PDF Pdf Storage And Clustering Enterprise Hat Red Ex436

Thank you unconditionally much for downloading **Pdf Storage And Clustering Enterprise Hat Red Ex436**. Most likely you have knowledge that, people have seen numerous times for their favorite books taking into account this Pdf Storage And Clustering Enterprise Hat Red Ex436, but stop in the works in harmful downloads.

Rather than enjoying a fine ebook subsequent to a mug of coffee in the afternoon, otherwise they juggled taking into consideration some harmful virus inside their computer. **Pdf Storage And Clustering Enterprise Hat Red Ex436** is understandable in our digital library an online permission to it is set as public hence you can download it instantly. Our digital library saves in merged countries, allowing you to get the most less latency times to download any of our books afterward this one. Merely said, the Pdf Storage And Clustering Enterprise Hat Red Ex436 is universally compatible gone any devices to read.

KEY=EX436 - CUNNINGHAM JAIDYN

IBM SYSTEM STORAGE TS7600 WITH PROTECTIER VERSION 3.3

IBM Redbooks This IBM® Redbooks® publication describes the IBM solution for data deduplication, the IBM System Storage® TS7650G IBM ProtecTIER® Deduplication Gateway, and the IBM TS7620 ProtecTIER Deduplication Appliance Express. This solution consists of the IBM System Storage ProtecTIER Enterprise Edition V3.3 software and the IBM System Storage TS7600 family of products. They are designed to address the disk-based data protection needs of enterprise data centers. We describe the components that make up IBM System Storage TS7600 with ProtecTIER and provide extensive planning and sizing guidance that enables you to determine your requirements and the correct configuration for your environment. We then guide you through the basic setup steps on the system and on the host. We also describe all operational tasks that are required during normal day-to-day operation or when upgrading your TS7600 products. All available models of the ProtecTIER deduplication system can now be ordered in a configuration to operate in one of the following modes for which we provide setup, configuration and usage guidelines for your business needs: The Virtual Tape Library (VTL) interface is the foundation of ProtecTIER and emulates traditional automated tape libraries. The Symantec NetBackup OpenStorage (OST) API can be integrated with Symantec NetBackup to provide backup-to-disk without having to emulate traditional tape libraries. The newly available File System Interface (FSI) supports Common Internet File System (CIFS) and Network File System (NFS) as a backup target. This publication is intended for system programmers, storage administrators, hardware and software planners, and other IT personnel that are involved in planning, implementing, and the use of the IBM deduplication solution. It also is intended for anyone seeking detailed technical information about the IBM System Storage TS7600 with ProtecTIER.

RED HAT OPENSIFT V4.3 ON IBM POWER SYSTEMS REFERENCE GUIDE

IBM Redbooks This IBM® Redpaper publication describes how to deploy Red Hat OpenShift V4.3 on IBM Power Systems servers. This book presents reference architectures for deployment, initial sizing guidelines for server, storage, and IBM Cloud® Paks. Moreover, this publication delivers information about initial supported Power System configurations for Red Hat OpenShift V4.3 deployment (bare metal, IBM PowerVM® LE LPARs, and others). This book serves as a guide for how to deploy Red Hat OpenShift V4.3 and provide start guidelines and recommended practices for implementing it on Power Systems and completing it with the supported IBM Cloud Paks. The publication addresses topics for developers, IT architects, IT specialists, sellers, and anyone who wants to implement a Red Hat OpenShift V4.3 and IBM Cloud Paks on IBM Power Systems. This book also provides technical content to transfer how-to skills to the support teams, and solution guidance to the sales team. This book complements the documentation that is available at IBM Knowledge Center, and also aligns with the educational offerings that are provided by the IBM Systems Technical Education (SSE).

IBM POWER SYSTEMS HIGH AVAILABILITY AND DISASTER RECOVERY UPDATES: PLANNING FOR A MULTICLOUD ENVIRONMENT

IBM Redbooks This IBM® Redpaper publication delivers an updated guide for high availability and disaster recovery (HADR) planning in a multicloud environment for IBM Power. This publication describes the ideas from studies that were performed in a virtual collaborative team of IBM Business Partners, technical focal points, and product managers who used hands-on experience to implement case studies to show HADR management aspects to develop this technical update guide for a hybrid multicloud environment. The goal of this book is to deliver a HADR guide for backup and data management on-premises and in a multicloud environment. This document updates HADR on-premises and in the cloud with IBM PowerHA® SystemMirror®, IBM VM Recovery Manager (VMRM), and other solutions that are available on IBM Power for IBM AIX®, IBM i, and Linux. This publication highlights the available offerings at the time of writing for each operating system (OS) that is supported in IBM Power, including best practices. This book addresses topics for IT architects, IT specialists, sellers, and anyone looking to implement and manage HADR on-premises and in the cloud. Moreover, this publication provides documentation to transfer how-to skills to the technical teams and solution guidance to the sales team. This book complements the documentation that is available at IBM Documentation and aligns with the educational materials that are provided by IBM Systems Technical Training.

GETTING STARTED WITH RED HAT ENTERPRISE VIRTUALIZATION

Packt Publishing Ltd If you are a system administrator who is interested in implementing and managing open source virtualization infrastructures, this is the book for you. A basic knowledge of virtualization and basic Linux command line experience is needed.

LINUX BIBLE

BOOT UP TO UBUNTU, FEDORA, KNOPPIX, DEBIAN, OPENSUSE, AND 11 OTHER DISTRIBUTIONS

John Wiley & Sons Demonstrates new Linux distributions while covering commands, installation, customizing the Linux shell, filesystem management, working with multimedia features, security, networking, and system administration.

LINUX BIBLE 2009 EDITION

BOOT UP UBUNTU, FEDORA, KNOPPIX, DEBIAN, OPENSUSE, AND MORE

John Wiley & Sons As the ultimate resource on the basics of the Linux desktops, servers, and programming tools, this book is written by the ultimate author on all things Linux. This transition resource is ideal if you are making the move from Windows or Macintosh to using Linux as your desktop system, and explains the Linux technology, offers detailed installation instructions, and includes step-by-step descriptions of key desktop and server components. You'll relish the in-depth descriptions that will help you choose the best Linux distribution to suit your needs. Note: CD-ROM/DVD and other supplementary materials are not included as part of eBook file.

IBM POWERVM VIRTUALIZATION INTRODUCTION AND CONFIGURATION

IBM Redbooks This IBM® Redbooks® publication provides an introduction to PowerVM™ 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.

SONAS IMPLEMENTATION AND BEST PRACTICES GUIDE

IBM Redbooks IBM® Scale Out Network Attached Storage (SONAS) is a Scale Out NAS offering designed to manage vast repositories of information in enterprise environments requiring very large capacities, high levels of performance, and high availability. The IBM SONAS appliance provides a range of reliable, scalable storage solutions for a variety of storage requirements. These capabilities are achieved by using network access protocols such as NFS, CIFS, HTTPS, FTP, and SCP. Using built-in RAID technologies, all data is well protected with options to add additional protection through mirroring, replication, snapshots, and backup. These storage systems are also characterized by simple management interfaces that make their installation, administration, and troubleshooting uncomplicated and straightforward. This IBM Redbooks® publication is the companion to the IBM Redbooks publication, SONAS Concepts, Architecture, and Planning Guide, SG24-7963. It is intended for storage administrators who have ordered their SONAS solution and are ready to install, customize, and use it. A quick start scenario takes you through common SONAS administration tasks to familiarize you with the SONAS system through the GUI and CLI. Backup and availability scenarios as well as best practices for setting up and troubleshooting hints and tips are included.

ACHIEVING HIGH AVAILABILITY ON LINUX FOR SYSTEM Z WITH LINUX-HA RELEASE 2

IBM Redbooks As Linux® on System z® becomes more prevalent and mainstream in the industry, the need for it to deliver higher levels of availability is increasing. IBM® supports the High Availability Linux (Linux-HA) project, which provides high availability functions to the open source community. One component of the Linux-HA project is the Heartbeat program, which runs on every known Linux platform. Heartbeat is part of the framework of the Linux-HA project. This IBM Redbooks® publication provides information to help you evaluate and implement Linux-HA release 2 by using Heartbeat 2.0 on the IBM System z platform with either SUSE® Linux Enterprise Server version 10 or Red Hat® Enterprise Linux® 5. To begin, we review the fundamentals of high availability concepts and terminology. Then we discuss the Heartbeat 2.0 architecture and its components. We examine some of the special considerations when using Heartbeat 2.0 on Linux on System z, particularly Linux on z/VM®, with logical partitions (LPARs), interguest communication by using HiperSockets™, and Shoot The Other Node In The Head (STONITH) by using VSMERVE for Simple Network IPL (snIPL). By reading this book, you can examine our environment as we outline our installation and setup processes and configuration. We demonstrate an active and passive single resource scenario and a quorum scenario by using a single resource with three guests in the cluster. Finally, we demonstrate and describe sample usage scenarios.

IBM POWER E1080 TECHNICAL OVERVIEW AND INTRODUCTION

IBM Redbooks This IBM® Redpaper® publication provides a broad understanding of a new architecture of the IBM Power® E1080 (also known as the Power E1080) server that supports IBM AIX®, IBM i, and selected distributions of Linux operating systems. The objective of this paper is to introduce the Power E1080, the most powerful and scalable server of the IBM Power portfolio, and its offerings and relevant functions: Designed to support up to four system nodes and up to 240 IBM Power10™ processor cores The Power E1080 can be initially ordered with a single system node or two system nodes configuration, which provides up to 60 Power10 processor cores with a single node configuration or up to 120 Power10 processor cores with a two system nodes configuration. More support for a three or four system nodes configuration is to be added on December 10, 2021, which provides support for up to 240 Power10 processor cores with a full combined four system nodes server. Designed to supports up to 64 TB memory The Power E1080 can be initially ordered with the total memory RAM capacity up to 8 TB. More support is to be added on December 10, 2021 to support up to 64 TB in a full combined four system nodes server. Designed to support up to 32 Peripheral Component Interconnect® (PCIe) Gen 5 slots in a full combined four system nodes server and up to 192 PCIe Gen 3 slots with expansion I/O drawers The Power E1080

supports initially a maximum of two system nodes; therefore, up to 16 PCIe Gen 5 slots, and up to 96 PCIe Gen 3 slots with expansion I/O drawer. More support is to be added on December 10, 2021, to support up to 192 PCIe Gen 3 slots with expansion I/O drawers. Up to over 4,000 directly attached serial-attached SCSI (SAS) disks or solid-state drives (SSDs) Up to 1,000 virtual machines (VMs) with logical partitions (LPARs) per system System control unit, providing redundant system master Flexible Service Processor (FSP) Supports IBM Power System Private Cloud Solution with Dynamic Capacity This publication is for professionals who want to acquire a better understanding of Power servers. The intended audience includes the following roles: Customers Sales and marketing professionals Technical support professionals IBM Business Partners Independent software vendors (ISVs) This paper does not replace the current marketing materials and configuration tools. It is intended as an extra source of information that, together with existing sources, can be used to enhance your knowledge of IBM server solutions.

SOFTWARE DEFINED DATA CENTER WITH RED HAT CLOUD AND OPEN SOURCE IT OPERATIONS MANAGEMENT

IBM Redbooks This IBM® Redbooks® publication delivers a Site Reliability Engineering (SRE) solution for cloud workloads that uses Red Hat OpenStack for Infrastructure as a Service (IaaS), Red Hat OpenShift for Platform as a Service (PaaS), and IT operations management that uses open source tools. Today, customers are no longer living in a world of licensed software. Curiosity increased the demand for investigating the Open Source world for Community Open Source and Enterprise grade applications. IBM as one of the contributors to the Open Source community is interested in helping the software be maintained and supported. Having companies, such as IBM, support the evolution of Open Source software helps to keep the Open Source community striving for enterprise grade open source solutions. Lately, companies are working on deciphering how to take advantage of Enterprise and Community Open Source to implement in their enterprises. The business case for open source software is no longer a mystery and no surprise that most of the new positions in IT enterprises are related to open source projects. The ability of a large enterprise to manage this sort of implementations is to engage in a hypertrophied cooperation, where the ability to not only cooperate with teams and people outside your organization, but also to find new ways of working together and devise new ways to improve the software and its code. A goal for this publication is to help the client's journey into the open source space and implement a private Cloud Container-based architecture with the ability to manage the entire IT Service Management processes from the open source framework. This publication describes the architecture and implementation details of the solution. Although not every piece of this solution is documented here, this book does provide instructions for what was achieved incorporating open source technologies. Moreover, with this publication, the team shares their collaboration experiences working in a team of technologists, open source developers, Red Hat, and the open source community. This publication is for designers, developers, managers, and anyone who is considering starting a Cloud open source project, or users who started that journey. This book also can be a manual to guide the implementation of a technical viable architecture and help those enterprises participate in an open source project but have not done so before. The reader must be familiar with principles in programming and basic software engineering concepts, such as source code, compilers, and patches.

IBM FLASHSYSTEM 5200 PRODUCT GUIDE

IBM Redbooks This IBM® Redbooks® Product Guide publication describes the IBM FlashSystem® 5200 solution, which is a next-generation IBM FlashSystem control enclosure. It is an NVMe end-to-end platform that is targeted at the entry and midrange market and delivers the full capabilities of IBM FlashCore® technology. It also provides a rich set of software-defined storage (SDS) features that are delivered by IBM Spectrum® Virtualize, including the following features: Data reduction and deduplication Dynamic tiering Thin provisioning Snapshots Cloning Replication Data copy services Transparent Cloud Tiering IBM HyperSwap® including 3-site replication for high availability (HA) Scale-out and scale-up configurations further enhance capacity and throughput for better availability. The IBM FlashSystem 5200 is a high-performance storage solution that is based on a revolutionary 1U form factor. It consists of 12 NVMe Flash Devices in a 1U storage enclosure drawer with full redundant canister components and no single point of failure. It is designed for businesses of all sizes, including small, remote, branch offices and regional clients. It is a smarter, self-optimizing solution that requires less management, which enables organizations to overcome their storage challenges. Flash has come of age and price point reductions mean that lower parts of the storage market are seeing the value of moving over to flash and NVMe--based solutions. The IBM FlashSystem 5200 advances this transition by providing incredibly dense tiers of flash in a more affordable package. With the benefit of IBM FlashCore Module compression and new QLC flash-based technology becoming available, a compelling argument exists to move away from Nearline SAS storage and on to NVMe. With the release of IBM FlashSystem 5200 Software V8.4, extra functions and features are available, including support for new Distributed RAID1 (DRAID1) features, GUI enhancements, Redirect-on-write for Data Reduction Pool (DRP) snapshots, and 3-site replication capabilities. This book is aimed at pre-sales and post-sales technical support and marketing and storage administrators.

DEVOPS WITH OPENSIFT

CLOUD DEPLOYMENTS MADE EASY

"O'Reilly Media, Inc." For many organizations, a big part of DevOps' appeal is software automation using infrastructure-as-code techniques. This book presents developers, architects, and infra-ops engineers with a more practical option. You'll learn how a container-centric approach from OpenShift, Red Hat's cloud-based PaaS, can help your team deliver quality software through a self-service view of IT infrastructure. Three OpenShift experts at Red Hat explain how to configure Docker application containers and the Kubernetes cluster manager with OpenShift's developer- and operational-centric tools. Discover how this infrastructure-agnostic container management platform can help companies navigate the murky area where infrastructure-as-code ends and application automation begins. Get an application-centric view of automation--and understand why it's important Learn patterns and practical examples for managing continuous deployments such as rolling, A/B, blue-green, and canary Implement continuous integration pipelines with OpenShift's Jenkins capability Explore mechanisms for separating and managing configuration from static runtime software Learn how to use and customize OpenShift's source-to-image capability Delve into management and operational

considerations when working with OpenShift-based application workloads Install a self-contained local version of the OpenShift environment on your computer

VERFAHREN UND WERKZEUGE ZUR LEISTUNGSMESSUNG, -ANALYSE UND -BEWERTUNG DER EIN-, AUSGABEEINHEITEN VON RECHENSYSTEMEN

Logos Verlag Berlin GmbH Die Leistungsfähigkeit moderner Rechensysteme erhöht sich stetig. Allerdings zeigen Untersuchungen, dass die Rechenleistung von Prozessoren starker steigt, als die Ein-/Ausgabeleistung zu den Sekundarspeichern, auf denen die Daten zur Berechnung abgelegt werden. Dies führt dazu, dass Prozessoren zukünftig nicht ihr gesamtes Rechenpotential ausschöpfen können, da sie auf Daten der Sekundarspeicher warten müssen. Damit die Leistung der Speicher nicht zu einem leistungsbegrenzenden Faktor des gesamten Systems wird, ist die Leistungsanalyse und -optimierung der Sekundarspeicher notwendig. Die Leistungsmessung von Sekundarspeichersystemen wird typischerweise mit Softwarewerkzeugen durchgeführt, die eine Last auf dem Sekundarspeicher erzeugen und anhand dieser Last eine Leistungsermittlung durchführen. Diese sogenannten I/O-Benchmarks haben zahlreiche Probleme, die im Rahmen dieser Dissertation aufgezeigt und gelöst werden. Es wird ein neuer Ansatz entwickelt, der realitätsnahes, nutzerrelevantes, vergleichbares und dennoch einfaches I/O-Benchmarking insbesondere in Hinblick auf die Leistungsermittlung beim Zugriff auf Sekundarspeicher mittels der MPI-IO-Schnittstelle ermöglicht. Ausgehend von den notwendigen Schritten bei der Leistungsanalyse wird eine neue Benchmark-Architektur entwickelt, die insbesondere Lösungen für die gefundenen Probleme der geringen Repräsentativität von Benchmarkergebnissen und der fehlenden Nutzerunterstützung beim Benchmarking bietet und damit über vorhandene Arbeiten in diesem Themenbereich deutlich hinausgeht. Es wird ein Benchmark-System erstellt, das nutzerrelevante Ergebnisse ermittelt, indem es dem Nutzer ermöglicht, das Lastverhalten MPI-IO-basierter Applikationen als Messgrundlage zu verwenden. Ausserdem wird eine realitätsnahe und einfach nutzbare I/O-Lastbeschreibung präsentiert, deren Möglichkeiten existierende I/O-Lastbeschreibungen in Hinblick auf Genauigkeit bei Verwendung komplexer paralleler I/O-Lasten übersteigen. Die Funktionsfähigkeit und Genauigkeit des I/O-Benchmarking-Ansatzes wird mit Messungen anhand von Beispielapplikationen gezeigt.

VMWARE COOKBOOK

A REAL-WORLD GUIDE TO EFFECTIVE VMWARE USE

"O'Reilly Media, Inc." Provides information on working with VMware ESXi in a variety of network environments, covering such topics as storage, networking, security, and vCloud.

KUBERNETES OPERATORS

AUTOMATING THE CONTAINER ORCHESTRATION PLATFORM

O'Reilly Media Operators are a way of packaging, deploying, and managing Kubernetes applications. A Kubernetes application doesn't just run on Kubernetes; it's composed and managed in Kubernetes terms. Operators add application-specific operational knowledge to a Kubernetes cluster, making it easier to automate complex, stateful applications and to augment the platform. Operators can coordinate application upgrades seamlessly, react to failures automatically, and streamline repetitive maintenance like backups. Think of Operators as site reliability engineers in software. They work by extending the Kubernetes control plane and API, helping systems integrators, cluster administrators, and application developers reliably deploy and manage key services and components. Using real-world examples, authors Jason Dobies and Joshua Wood demonstrate how to use Operators today and how to create Operators for your applications with the Operator Framework and SDK. Learn how to establish a Kubernetes cluster and deploy an Operator Examine a range of Operators from usage to implementation Explore the three pillars of the Operator Framework: the Operator SDK, the Operator Lifecycle Manager, and Operator Metering Build Operators from the ground up using the Operator SDK Build, package, and run an Operator in development, testing, and production phases Learn how to distribute your Operator for installation on Kubernetes clusters

OPENSIFT IN ACTION

Simon and Schuster Summary OpenShift in Action is a full reference to Red Hat OpenShift that breaks down this robust container platform so you can use it day-to-day. Combining Docker and Kubernetes, OpenShift is a powerful platform for cluster management, scaling, and upgrading your enterprise apps. It doesn't matter why you use OpenShift—by the end of this book you'll be able to handle every aspect of it, inside and out! Foreword by Jim Whitehurst, Red Hat. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Containers let you package everything into one neat place, and with Red Hat OpenShift you can build, deploy, and run those packages all in one place! Combining Docker and Kubernetes, OpenShift is a powerful platform for cluster management, scaling, and upgrading your enterprise apps. About the Book OpenShift in Action is a full reference to Red Hat OpenShift that breaks down this robust container platform so you can use it day-to-day. Starting with how to deploy and run your first application, you'll go deep into OpenShift. You'll discover crystal-clear explanations of namespaces, cgroups, and SELinux, learn to prepare a cluster, and even tackle advanced details like software-defined networks and security, with real-world examples you can take to your own work. It doesn't matter why you use OpenShift—by the end of this book you'll be able to handle every aspect of it, inside and out! What's Inside Written by lead OpenShift architects Rock-solid fundamentals of Docker and Kubernetes Keep mission-critical applications up and running Manage persistent storage About the Reader For DevOps engineers and administrators working in a Linux-based distributed environment. About the Authors Jamie Duncan is a cloud solutions architect for Red Hat, focusing on large-scale OpenShift deployments. John Osborne is a principal OpenShift architect for Red Hat. Table of Contents PART 1 - FUNDAMENTALS Getting to know OpenShift Getting started Containers are Linux PART 2 - CLOUD-NATIVE APPLICATIONS Working with services Autoscaling with metrics Continuous integration and continuous deployment PART 3 - STATEFUL

APPLICATIONS Creating and managing persistent storage Stateful applications PART 4 - OPERATIONS AND SECURITY Authentication and resource access Networking Security

IMPLEMENTING THE IBM STORWIZE

IBM Redbooks Organizations of all sizes are faced with the challenge of managing massive volumes of increasingly valuable data. However, storing this data can be costly, and extracting value from the data is becoming more and more difficult. IT organizations have limited resources, but must stay responsive to dynamic environments and act quickly to consolidate, simplify, and optimize their IT infrastructures. The IBM® Storwize® V3700 system provides a solution that is affordable, easy to use, and self-optimizing, which enables organizations to overcome these storage challenges. Storwize V3700 delivers efficient, entry-level configurations that are specifically designed to meet the needs of small and midsize businesses. Designed to provide organizations with the ability to consolidate and share data at an affordable price, Storwize V3700 offers advanced software capabilities that are usually found in more expensive systems. Built on innovative IBM technology, Storwize V3700 addresses the block storage requirements of small and midsize organizations, Storwize V3700 is designed to accommodate the most common storage network technologies. This design enables easy implementation and management. Storwize V3700 includes the following features: Web-based GUI provides point-and-click management capabilities. Internal disk storage virtualization enables rapid, flexible provisioning and simple configuration changes. Thin provisioning enables applications to grow dynamically, but only use space they actually need. Enables simple data migration from external storage to Storwize V3700 storage (one-way from another storage device). Remote Mirror creates copies of data at remote locations for disaster recovery. IBM FlashCopy® creates instant application copies for backup or application testing. This IBM Redbooks® publication is intended for pre-sales and post-sales technical support professionals and storage administrators. The concepts in this book also relate to the IBM Storwize V3500. This book was written at a software level of version 7 release 4.

IBM FLASHSYSTEM 7200 PRODUCT GUIDE

IBM Redbooks This IBM® Redbooks® Product Guide publication describes the IBM FlashSystem® 7200 solution, which is a comprehensive, all-flash, and NVMe-enabled enterprise storage solution that delivers the full capabilities of IBM FlashCore® technology. In addition, it provides a rich set of software-defined storage (SDS) features, including data reduction and de-duplication, dynamic tiering, thin-provisioning, snapshots, cloning, replication, data copy services, and IBM HyperSwap® for high availability (HA). Scale-out and scale-up configurations further enhance capacity and throughput for better availability

IBM SYSTEM STORAGE SAN VOLUME CONTROLLER, IBM STORWIZE V7000, AND IBM FLASHSYSTEM 7200 BEST PRACTICES AND PERFORMANCE GUIDELINES

IBM Redbooks This IBM® Redbooks® publication captures several of the preferred practices and describes the performance gains that can be achieved by implementing the IBM System Storage® SAN Volume Controller and IBM Storwize® V7000 powered by IBM Spectrum Virtualize™ V8.2.1. These practices are based on field experience. This book highlights configuration guidelines and preferred practices for the storage area network (SAN) topology, clustered system, back-end storage, storage pools and managed disks, volumes, remote copy services, and hosts. Then it provides performance guidelines for SAN Volume Controller, back-end storage, and applications. It explains how you can optimize disk performance with the IBM System Storage Easy Tier® function. It also provides preferred practices for monitoring, maintaining, and troubleshooting SAN Volume Controller and Storwize V7000. This book is intended for experienced storage, SAN, and SAN Volume Controller administrators and technicians. Understanding this book requires advanced knowledge of the SAN Volume Controller and Storwize V7000 and SAN environments. Important: On 11th February 2020 IBM announced the arrival of SAN Volume Controller SA2 and SV2, and IBM FlashSystem® 7200 to the family. This book was written specifically for prior versions of SVC and Storwize V7000; however, most of the general principles will apply. If you are in any doubt as to their applicability then you should work with your local IBM representative. This book will be updated to comprehensively include SAN Volume Controller SA2 and SV2 and FlashSystem 7200 in due course.

LINUX CLUSTERING

BUILDING AND MAINTAINING LINUX CLUSTERS

Sams Publishing "Linux Clustering" is the premier resource for system administrators wishing to implement clustering solutions on the many types of Linux systems. It guides Linux Administrators through difficult tasks while offering helpful tips and tricks.

POWERHA SYSTEMMIRROR FOR IBM I COOKBOOK

IBM Redbooks IBM® PowerHATM SystemMirror for i is the IBM high-availability disk-based clustering solution for the IBM i 7.1 operating system. When combined with IBM i clustering technology, PowerHA for i delivers a complete high-availability and disaster-recovery solution for your business applications running in the IBM System i® environment. PowerHA for i enables you to support high-availability capabilities with either native disk storage or IBM DS8000® or DS6000TM storage servers or IBM Storwize V7000 and SAN Volume Controllers. The latest release of IBM PowerHA SystemMirror for i delivers a brand-new web-based PowerHA graphical user interface that effectively combines the solution-based and task-based activities for your HA environment, all in a single user interface. This IBM Redbooks® publication provides a broad understanding of PowerHA for i. This book is intended for all IBM i professionals who are planning on implementing a PowerHA solution on IBM i.

IBM FLASHSYSTEM 9200 PRODUCT GUIDE

IBM Redbooks This IBM® Redbooks® Product Guide publication describes the IBM FlashSystem® 9200 solution, which is a comprehensive, all-flash, and NVMe-enabled enterprise storage solution that delivers the full capabilities of IBM FlashCore® technology. In addition, it provides a rich set of software-defined storage (SDS) features, including data reduction and de-duplication,

dynamic tiering, thin-provisioning, snapshots, cloning, replication, data copy services, and IBM HyperSwap® for high availability (HA). Scale-out and scale-up configurations further enhance capacity and throughput for better availability.

LINUX JOURNAL

97 THINGS EVERY CLOUD ENGINEER SHOULD KNOW

O'Reilly Media If you create, manage, operate, or configure systems running in the cloud, you're a cloud engineer--even if you work as a system administrator, software developer, data scientist, or site reliability engineer. With this book, professionals from around the world provide valuable insight into today's cloud engineering role. These concise articles explore the entire cloud computing experience, including fundamentals, architecture, and migration. You'll delve into security and compliance, operations and reliability, and software development. And examine networking, organizational culture, and more. You're sure to find 1, 2, or 97 things that inspire you to dig deeper and expand your own career. "Three Keys to Making the Right Multicloud Decisions," Brendan O'Leary "Serverless Bad Practices," Manases Jesus Galindo Bello "Failing a Cloud Migration," Lee Atchison "Treat Your Cloud Environment as If It Were On Premises," Iyana Garry "What Is Toil, and Why Are SREs Obsessed with It?," Zachary Nickens "Lean QA: The QA Evolving in the DevOps World," Theresa Neate "How Economies of Scale Work in the Cloud," Jon Moore "The Cloud Is Not About the Cloud," Ken Corless "Data Gravity: The Importance of Data Management in the Cloud," Geoff Hughes "Even in the Cloud, the Network Is the Foundation," David Murray "Cloud Engineering Is About Culture, Not Containers," Holly Cummins

KUBERNETES PATTERNS

REUSABLE ELEMENTS FOR DESIGNING CLOUD-NATIVE APPLICATIONS

O'Reilly Media The way developers design, build, and run software has changed significantly with the evolution of microservices and containers. These modern architectures use new primitives that require a different set of practices than most developers, tech leads, and architects are accustomed to. With this focused guide, Bilgin Ibryam and Roland Huß from Red Hat provide common reusable elements, patterns, principles, and practices for designing and implementing cloud-native applications on Kubernetes. Each pattern includes a description of the problem and a proposed solution with Kubernetes specifics. Many patterns are also backed by concrete code examples. This book is ideal for developers already familiar with basic Kubernetes concepts who want to learn common cloud native patterns. You'll learn about the following pattern categories: Foundational patterns cover the core principles and practices for building container-based cloud-native applications. Behavioral patterns explore finer-grained concepts for managing various types of container and platform interactions. Structural patterns help you organize containers within a pod, the atom of the Kubernetes platform. Configuration patterns provide insight into how application configurations can be handled in Kubernetes. Advanced patterns covers more advanced topics such as extending the platform with operators.

IMPLEMENTING THE IBM STORWIZE V3500

IBM Redbooks 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.

STRENGTHENING FORENSIC SCIENCE IN THE UNITED STATES

A PATH FORWARD

National Academies Press Scores of talented and dedicated people serve the forensic science community, performing vitally important work. However, they are often constrained by lack of adequate resources, sound policies, and national support. It is clear that change and advancements, both systematic and scientific, are needed in a number of forensic science disciplines to ensure the reliability of work, establish enforceable standards, and promote best practices with consistent application. Strengthening Forensic Science in the United States: A Path Forward provides a detailed plan for addressing these needs and suggests the creation of a new government entity, the National Institute of Forensic Science, to establish and enforce standards within the forensic science community. The benefits of improving and regulating the forensic science disciplines are clear: assisting law enforcement officials, enhancing homeland security, and reducing the risk of wrongful conviction and exoneration. Strengthening Forensic Science in the United States gives a full account of what is needed to advance the forensic science disciplines, including upgrading of systems and organizational structures, better training, widespread adoption of uniform and enforceable best practices, and mandatory certification and accreditation programs. While this book provides an essential call-to-action for congress and policy makers, it also serves as a vital tool for law enforcement agencies, criminal prosecutors and attorneys, and forensic science educators.

IBM SAN VOLUME CONTROLLER BEST PRACTICES AND PERFORMANCE GUIDELINES

IBM Redbooks This IBM® Redbooks® publication describes several of the preferred practices and describes the performance gains that can be achieved by implementing the IBM SAN Volume Controller powered by IBM Spectrum® Virtualize V8.4. These practices are based on field experience. This book highlights configuration guidelines and preferred practices for the storage area network (SAN)

topology, clustered system, back-end storage, storage pools, and managed disks, volumes, Remote Copy services, and hosts. Then, it provides performance guidelines for IBM SAN Volume Controller, back-end storage, and applications. It explains how you can optimize disk performance with the IBM System Storage Easy Tier® function. It also provides preferred practices for monitoring, maintaining, and troubleshooting IBM SAN Volume Controller. This book is intended for experienced storage, SAN, and IBM SAN Volume Controller administrators and technicians. Understanding this book requires advanced knowledge of the IBM SAN Volume Controller, IBM FlashSystem, and SAN environments.

IBM SOFTWARE DEFINED ENVIRONMENT

IBM Redbooks This IBM® Redbooks® publication introduces the IBM Software Defined Environment (SDE) solution, which helps to optimize the entire computing infrastructure--compute, storage, and network resources--so that it can adapt to the type of work required. In today's environment, resources are assigned manually to workloads, but that happens automatically in a SDE. In an SDE, workloads are dynamically assigned to IT resources based on application characteristics, best-available resources, and service level policies so that they deliver continuous, dynamic optimization and reconfiguration to address infrastructure issues. Underlying all of this are policy-based compliance checks and updates in a centrally managed environment. Readers get a broad introduction to the new architecture. Think integration, automation, and optimization. Those are enablers of cloud delivery and analytics. SDE can accelerate business success by matching workloads and resources so that you have a responsive, adaptive environment. With the IBM Software Defined Environment, infrastructure is fully programmable to rapidly deploy workloads on optimal resources and to instantly respond to changing business demands. This information is intended for IBM sales representatives, IBM software architects, IBM Systems Technology Group brand specialists, distributors, resellers, and anyone who is developing or implementing SDE.

BUILDING CLUSTERED LINUX SYSTEMS

Prentice Hall Until now, building and managing Linux clusters has required more intimate and specialized knowledge than most IT organizations possess. This book dramatically lowers the learning curve, bringing together all the hands-on knowledge and step-by-step techniques needed to get the job done.

KUBERNETES: UP AND RUNNING

DIVE INTO THE FUTURE OF INFRASTRUCTURE

"O'Reilly Media, Inc." Legend has it that Google deploys over two billion application containers a week. How's that possible? Google revealed the secret through a project called Kubernetes, an open source cluster orchestrator (based on its internal Borg system) that radically simplifies the task of building, deploying, and maintaining scalable distributed systems in the cloud. This practical guide shows you how Kubernetes and container technology can help you achieve new levels of velocity, agility, reliability, and efficiency. Authors Kelsey Hightower, Brendan Burns, and Joe Beda—who've worked on Kubernetes at Google and other organizations—explain how this system fits into the lifecycle of a distributed application. You will learn how to use tools and APIs to automate scalable distributed systems, whether it is for online services, machine-learning applications, or a cluster of Raspberry Pi computers. Explore the distributed system challenges that Kubernetes addresses Dive into containerized application development, using containers such as Docker Create and run containers on Kubernetes, using the docker image format and container runtime Explore specialized objects essential for running applications in production Reliably roll out new software versions without downtime or errors Get examples of how to develop and deploy real-world applications in Kubernetes

PROFESSIONAL RED HAT ENTERPRISE LINUX 3

John Wiley & Sons What is this book about? Professional Red Hat Enterprise Linux 3 is a complete professional guide to setting up, configuring, and deploying Red Hat Enterprise Linux in the corporate production environment. The book focuses on Enterprise Server and Advanced Server features, including the key areas of high availability with the Red Hat Cluster Suite, Red Hat Network Control Center, and Red Hat Enterprise applications such as the Content Management System and portal server. Other key unique features include kernel tuning for various performance profiles; advanced Apache configuration; Tux installation/maintenance; building high-performance FTP servers; building high-performance mail servers (which means replacing Sendmail); Mailing list management; how to efficiently add, remove, or modify 100 users at the same time; and a discussion of disk quota management and monitoring. What does this book cover? The key features of the book include the following: How to install and setup RHEL 3 How to deploy RHEL 3 in production environment How to manage an RHEL system using Perl and shell scripting Advanced administration tools How to use Red Hat network service Details on installation and setup of security tools Ability to use and deploy High Availability solutions provided with RHEL 3 Performance tuning How to use monitoring tools Ability to use RHEL to provide scalable infrastructure solutions.

IMPLEMENTING THE IBM STORWIZE V7000 GEN2

IBM Redbooks Data is the new currency of business, the most critical asset of the modern organization. In fact, enterprises that can gain business insights from their data are twice as likely to outperform their competitors. Nevertheless, 72% of them have not started, or are only planning, big data activities. In addition, organizations often spend too much money and time managing where their data is stored. The average firm purchases 24% more storage every year, but uses less than half of the capacity that it already has. The IBM® Storwize® family, including the IBM SAN Volume Controller Data Platform, is a storage virtualization system that enables a single point of control for storage resources. This functionality helps support improved business application availability and greater resource use. The following list describes the business objectives of this system: To manage storage resources in your information technology (IT) infrastructure To make sure that those resources are used to the advantage of your business To do it quickly, efficiently, and in real time, while avoiding increases in administrative costs Virtualizing storage with Storwize helps make new and existing storage more effective. Storwize includes many functions traditionally deployed separately in disk systems. By including these

functions in a virtualization system, Storwize standardizes them across virtualized storage for greater flexibility and potentially lower costs. Storwize functions benefit all virtualized storage. For example, IBM Easy Tier® optimizes use of flash memory. In addition, IBM Real-time Compression™ enhances efficiency even further by enabling the storage of up to five times as much active primary data in the same physical disk space. Finally, high-performance thin provisioning helps automate provisioning. These benefits can help extend the useful life of existing storage assets, reducing costs. Integrating these functions into Storwize also means that they are designed to operate smoothly together, reducing management effort. This IBM Redbooks® publication provides information about the latest features and functions of the Storwize V7000 Gen2 and software version 7.3 implementation, architectural improvements, and Easy Tier.

RED HAT ENTERPRISE LINUX 8 ADMINISTRATION

MASTER LINUX ADMINISTRATION SKILLS AND PREPARE FOR THE RHCSA CERTIFICATION EXAM

Packt Publishing Ltd Develop the skill to manage and administer Red Hat Enterprise Linux and get ready to achieve the RHCSA certification Key Features Learn the most common administration and security tasks and manage enterprise Linux infrastructures efficiently Assess your knowledge using self-assessment questions based on real-world examples Understand how to apply the concepts of core systems administration in the real world Book Description Whether in infrastructure or development, as a DevOps or site reliability engineer, Linux skills are now more relevant than ever for any IT job, forming the foundation of understanding the most basic layer of your architecture. With Red Hat Enterprise Linux (RHEL) becoming the most popular choice for enterprises worldwide, achieving the Red Hat Certified System Administrator (RHCSA) certification will validate your Linux skills to install, configure, and troubleshoot applications and services on RHEL systems. Complete with easy-to-follow tutorial-style content, self-assessment questions, tips, best practices, and practical exercises with detailed solutions, this book covers essential RHEL commands, user and group management, software management, networking fundamentals, and much more. You'll start by learning how to create an RHEL 8 virtual machine and get to grips with essential Linux commands. You'll then understand how to manage users and groups on an RHEL 8 system, install software packages, and configure your network interfaces and firewall. As you advance, the book will help you explore disk partitioning, LVM configuration, Stratis volumes, disk compression with VDO, and container management with Podman, Buildah, and Skopeo. By the end of this book, you'll have covered everything included in the RHCSA EX200 certification and be able to use this book as a handy, on-the-job desktop reference guide. This book and its contents are solely the work of Miguel Perez Colino, Pablo Iranzo Gomez, and Scott McCarty. The content does not reflect the views of their employer (Red Hat Inc.). This work has no connection to Red Hat, Inc. and is not endorsed or supported by Red Hat, Inc. What you will learn Deploy RHEL 8 in different footprints, from bare metal and virtualized to the cloud Manage users and software on local and remote systems at scale Discover how to secure a system with SELinux, OpenSCAP, and firewalld Gain an overview of storage components with LVM, Stratis, and VDO Master remote administration with passwordless SSH and tunnels Monitor your systems for resource usage and take actions to fix issues Understand the boot process, performance optimizations, and containers Who this book is for This book is for IT professionals or students who want to start a career in Linux administration and anyone who wants to take the RHCSA 8 certification exam. Basic knowledge of Linux and familiarity with the Linux command-line is necessary.

BUILDING STORAGE NETWORKS

McGraw-Hill/Osborne Media "Develop and design successful storage systems using this in-depth resource, now in a completely revised second edition. Covering everything from basic fundamentals - such as I/O components and file systems to emerging topics such as i-SCSI and DAFS - this book delivers the background information and technical know-how to implement large-capacity, high-availability storage networks throughout your enterprise. Filled with diagrams and easy-to-understand explanations, this book will help you identify and apply network storage technology to best meet the needs of your organization."--BOOK JACKET.Title Summary field provided by Blackwell North America, Inc. All Rights Reserved

DEPLOYING TO OPENSIFT

A GUIDE FOR BUSY DEVELOPERS

O'Reilly Media, Inc. "Get an in-depth tour of OpenShift, the container-based software deployment and management platform from Red Hat that provides a secure multi-tenant environment for the enterprise. This practical guide describes in detail how OpenShift, building on Kubernetes, enables you to automate the way you create, ship, and run applications in a containerized environment. Author Graham Dumpleton provides the knowledge you need to make the best use of the OpenShift container platform to deploy not only your cloud-native applications, but also more traditional stateful applications. Developers and administrators will learn how to run, access, and manage containers in OpenShift, including how to orchestrate them at scale. Build application container images from source and deploy them Implement and extend application image builders Use incremental and chained builds to accelerate build times Automate builds by using a webhook to link OpenShift to a Git repository Add configuration and secrets to the container as project resources Make an application visible outside the OpenShift cluster Manage persistent storage inside an OpenShift container Monitor application health and manage the application lifecycle This book is a perfect follow-up to OpenShift for Developers: A Guide for Impatient Beginners (O'Reilly).

LOTUS DOMINO R5 CLUSTERING

Prentice Hall In this book, IBM and Lotus insiders demonstrate exactly how to make the most of Domino 5's exceptionally powerful clusters to maximize both availability and scalability. The authors cover every phase of clustering: planning and design, installation and configuration, day-to-day management, troubleshooting, and more. Learn how to use Domino R5's new Internet Cluster Manager (ICM) to extend failover and load balancing capabilities to Web browsers, enhancing the reliability of any Web applications. Coverage

also includes: clustering multiple servers in cross-platform environments; clustering combinations of R5 and R4.5 servers; deploying multiple clusters in a Domino domain; including partitioned servers within clusters; and spanning multiple LAN segments to create campus-wide clusters. Many examples utilize Windows 2000, but this is the only Lotus Domino R5 clustering guide that also provides full coverage of clustering in Linux environments.

IBM SYSTEM STORAGE N SERIES CLUSTERED DATA ONTAP

IBM Redbooks IBM® System Storage® N series storage systems offer an excellent solution for a broad range of deployment scenarios. IBM System Storage N series storage systems function as a multiprotocol storage device that is designed to allow you to simultaneously serve both file and block-level data across a single network. These activities are demanding procedures that, for some solutions, require multiple, separately managed systems. The flexibility of IBM System Storage N series storage systems, however, allows them to address the storage needs of a wide range of organizations, including distributed enterprises and data centers for midrange enterprises. IBM System Storage N series storage systems also support sites with computer and data-intensive enterprise applications, such as database, data warehousing, workgroup collaboration, and messaging. This IBM Redbooks® publication explains the software features of the IBM System Storage N series storage systems with Clustered Data ONTAP (cDOT) Version 8.2, which is the first version available on the IBM System Storage N series, and as of October 2013, is also the most current version available. cDOT is different from previous ONTAP versions by the fact that it offers a storage solution that operates as a cluster with flexible scaling capabilities. cDOT configurations allow clients to build a scale-out architecture, protecting their investment and allowing horizontal scaling of their environment. This book also covers topics such as installation, setup, and administration of those software features from the IBM System Storage N series storage systems and clients, and provides example scenarios.

MASTERING CLOUDFORMS AUTOMATION

AN ESSENTIAL GUIDE FOR CLOUD ADMINISTRATORS

"O'Reilly Media, Inc." Learn how to work with the Automate feature of CloudForms, the powerful Red Hat cloud management platform that lets you administer your virtual infrastructure, including hybrid public and private clouds. This practical hands-on introduction shows you how to increase your operational efficiency by automating day-to-day tasks that now require manual input. Throughout the book, author Peter McGowan provides a combination of theoretical information and practical coding examples to help you learn the Automate object model. With this CloudForms feature, you can create auto-scalable cloud applications, eliminate manual decisions and operations when provisioning virtual machines and cloud instances, and manage your complete virtual machine lifecycle. In six parts, this book helps you: Learn the objects and concepts for developing automation scripts with CloudForms Automate Customize the steps and workflows involved in provisioning virtual machines Create and use service catalogs, items, dialogs, objects, bundles, and hierarchies Use CloudForm's updated workflow to retire and delete virtual machines and services Orchestrate and coordinate with external services as part of a workflow Explore distributed automation processing as well as argument passing and handling