

# Where To Download Clusters For High Availability A Primer Of Hp Ux Solutions Free Download Pdf

*High Availability MySQL High Availability Pro Linux High Availability Clustering High Availability and Disaster Recovery Blueprints for High Availability MySQL High Availability NGINX Cookbook Cassandra High Availability Clusters for High Availability Blueprints for High Availability Understanding and Using Q Replication for High Availability Solutions on the IBM z/OS Platform JUNOS High Availability Microsoft SQL Server 2000 High Availability End-to-End High Availability Solution for System z from a Linux Perspective CentOS High Availability Juniper SRX Series Pro SQL Server Always On Availability Groups High Availability and Disaster Recovery Options for DB2 for Linux, UNIX, and Windows Pro SharePoint Disaster Recovery and High Availability Professional SQL Server High Availability and Disaster Recovery Architecting for Scale SQL Server 2016 High Availability Unleashed (includes Content Update Program) IBM High Availability Solution for IBM FileNet P8 Systems High Availability and Scalability of Mainframe Environments Using System Z and Z/OS as Example High Availability MySQL Cookbook Designing High Availability Systems PostgreSQL High Availability Cookbook - Second Edition Linux in Action Architecting for Scale Oracle High Availability, Disaster Recovery, and Cloud Services PostgreSQL 12 High Availability Cookbook ServiceNow Application Development Microsoft Exchange Server 2013 Inside Out Mailbox and High Availability Apache Karaf Cookbook Teach Yourself Linux Virtualization and High Availability High Performance Drupal JBoss EAP6 High Availability High Availability Network Fundamentals Big Data Analytics with Hadoop 3 High Performance MySQL*

**Clusters for High Availability** Jun 11 2022 The expert guide to high availability clusters for HP-UX, Linux, Windows 2000, and Windows NT. The start-to-finish guide to high availability clustering Includes ways to maximize enterprise application availability—and minimize cost Completely updated for the latest tools, technologies, and applications Describes high availability solutions in HP-UX, Linux, and Windows environments Business-critical applications require higher availability than ever before—and today's high availability systems rely on clustering as a key strategy for maximizing reliability and robustness. In *Clusters for High Availability, Second Edition*, Peter S. Weygant covers all three pillars of successful high availability computing: robust technology, sound computing processes, and proactive support. He addresses every aspect of delivering high availability clustered systems: terminology, architecture, implementation, management,

monitoring, and beyond. Coverage includes: Fundamental concepts and components associated with high availability clustering A 14-step checklist for assessing your high availability needs Clustering techniques for HP-UX, Windows 2000, Windows NT, and Linux Clustered storage, backup, and network infrastructure solutions Practical techniques for building "disaster-tolerant" systems State-of-the-art cluster replication, monitoring, and management tools Weygant presents several brand-new case studies, including an Oracle Parallel Server application providing 5nines:5minutes protection; a high availability brokerage application built using a continental cluster; and a storage area network solution designed for an Internet service provider. The book also contains an extensive glossary. If you're responsible for delivering high availability, Clusters for High Availability is the comprehensive, up-to-date blueprint you need.

MySQL High Availability Sep 14 2022 Server bottlenecks and failures are a fact of life in any database deployment, but they don't have to bring everything to a halt. MySQL has several features that can help you protect your system from outages, whether it's running on hardware, virtual machines, or in the cloud. MySQL High Availability explains how to use these replication, cluster, and monitoring features in a wide range of real-life situations. Written by engineers who designed many of the tools covered inside, this book reveals undocumented or hard-to-find aspects of MySQL reliability and high availability -- knowledge that's essential for any organization using this database system. Explore the binary log, a file for replication that helps in disaster recovery and troubleshooting Get techniques for improving response time and handling large data sets Monitor database activity and performance, as well as major operating system parameters Keep track of what masters and slaves are doing, and deal with failures and restarts, corruption, and other incidents Automate key tasks with code from an open source library written by the authors Learn techniques for using MySQL in virtualized environments, such as Amazon Web Services Use MySQL Cluster to achieve high availability "MySQL replication is widely deployed but has never been adequately explained. This book changes that."-- Mark Callaghan, MySQL contributor and leader of MySQL engineering efforts at a few of the world's largest Internet companies

Linux in Action Oct 23 2020 Summary Linux in Action is a task-based tutorial that will give you the skills and deep understanding you need to administer a Linux-based system. This hands-on book guides you through 12 real-world projects so you can practice as you learn. Each chapter ends with a review of best practices, new terms, and exercises. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology You can't learn anything without getting your hands dirty--including Linux. Skills like securing files, folders, and servers, safely installing patches and applications, and managing a network are required for any serious user, including developers, administrators, and DevOps professionals. With this hands-on tutorial, you'll roll up your sleeves and learn Linux project by project. About the Book Linux in Action guides you through 12 real-world projects, including automating a backup-and-restore system, setting up a private Dropbox-style file cloud, and building your own MediaWiki server. You'll try out interesting examples as you lock in core practices like virtualization, disaster recovery, security, backup, DevOps, and system troubleshooting. Each chapter ends with a review of best practices, new terms, and

exercises. What's inside Setting up a safe Linux environment Managing secure remote connectivity Building a system recovery device Patching and upgrading your system About the Reader No prior Linux admin experience is required. About the Author David Clinton is a certified Linux Server Professional, seasoned instructor, and author of Manning's bestselling Learn Amazon Web Services in a Month of Lunches. Table of Contents Welcome to Linux Linux virtualization: Building a Linux working environment Remote connectivity: Safely accessing networked machines Archive management: Backing up or copying entire file systems Automated administration: Configuring automated offsite backups Emergency tools: Building a system recovery device Web servers: Building a MediaWiki server Networked file sharing: Building a Nextcloud file-sharing server Securing your web server Securing network connections: Creating a VPN or DMZ System monitoring: Working with log files Sharing data over a private network Troubleshooting system performance issues Troubleshooting network issues Troubleshooting peripheral devices DevOps tools: Deploying a scripted server environment using Ansible

PostgreSQL 12 High Availability Cookbook Jul 20 2020 A comprehensive guide to understanding key techniques for architecture and hardware planning, monitoring, replication, backups, and decoupling Key Features Newly updated edition, covering the latest PostgreSQL 12 features with hands-on industry-driven recipes Create a PostgreSQL cluster that stays online even when disaster strikes Learn how to avoid costly downtime and data loss that can ruin your business Book Description Databases are nothing without the data they store. In the event of an outage or technical catastrophe, immediate recovery is essential. This updated edition ensures that you will learn the important concepts related to node architecture design, as well as techniques such as using repmgr for failover automation. From cluster layout and hardware selection to software stacks and horizontal scalability, this PostgreSQL cookbook will help you build a PostgreSQL cluster that will survive crashes, resist data corruption, and grow smoothly with customer demand. You'll start by understanding how to plan a PostgreSQL database architecture that is resistant to outages and scalable, as it is the scaffolding on which everything rests. With the bedrock established, you'll cover the topics that PostgreSQL database administrators need to know to manage a highly available cluster. This includes configuration, troubleshooting, monitoring and alerting, backups through proxies, failover automation, and other considerations that are essential for a healthy PostgreSQL cluster. Later, you'll learn to use multi-master replication to maximize server availability. Later chapters will guide you through managing major version upgrades without downtime. By the end of this book, you'll have learned how to build an efficient and adaptive PostgreSQL 12 database cluster. What you will learn Understand how to protect data with PostgreSQL replication tools Focus on hardware planning to ensure that your database runs efficiently Reduce database resource contention with connection pooling Monitor and visualize cluster activity with Nagios and the TIG (Telegraf, InfluxDB, Grafana) stack Construct a robust software stack that can detect and avert outages Use multi-master to achieve an enduring PostgreSQL cluster Who this book is for This book is for Postgres administrators and developers who are looking to build and maintain a highly reliable PostgreSQL cluster. Although knowledge of the new features

of PostgreSQL 12 is not required, a basic understanding of PostgreSQL administration is expected.

**Understanding and Using Q Replication for High Availability Solutions on the IBM z/OS Platform** Apr 09 2022 With ever-increasing workloads on production systems from transaction, batch, online query and reporting applications, the challenges of high availability and workload balancing are more important than ever. This IBM® Redbooks® publication provides descriptions and scenarios for high availability solutions using the Q Replication technology of the IBM InfoSphere® Data Replication product on the IBM z/OS® platform. Also included are key considerations for designing, implementing, and managing solutions for the typical business scenarios that rely on Q Replication for their high availability solution. This publication also includes sections on latency analysis, managing Q Replication in the IBM DB2® for z/OS environment, and recovery procedures. These are topics of particular interest to clients who implement the Q Replication solution on the z/OS platform. Q Replication is a high-volume, low-latency replication solution that uses IBM WebSphere® MQ message queues to replicate transactions between source and target databases or subsystems. A major business benefit of the low latency and high throughput solution is timely availability of the data where the data is needed. High availability solutions are implemented to minimize the impact of planned and unplanned disruptions of service to the applications. Disruption of service can be caused by software maintenance and upgrades or by software and hardware outages. As applications' high availability requirements evolve towards continuous availability, that is availability of the data 24 hours a day and 7 days a week, so does the Q Replication solution, to meet these challenges. If you are interested in the Q Replication solution and how it can be used to implement some of the high availability requirements of your business scenarios, this book is for you.

**Architecting for Scale** Sep 21 2020 Every day, companies struggle to scale critical applications. As traffic volume and data demands increase, these applications become more complicated and brittle, exposing risks and compromising availability. This practical guide shows IT, devops, and system reliability managers how to prevent an application from becoming slow, inconsistent, or downright unavailable as it grows. Scaling isn't just about handling more users; it's also about managing risk and ensuring availability. Author Lee Atchison provides basic techniques for building applications that can handle huge quantities of traffic, data, and demand without affecting the quality your customers expect. In five parts, this book explores: Availability: learn techniques for building highly available applications, and for tracking and improving availability going forward Risk management: identify, mitigate, and manage risks in your application, test your recovery/disaster plans, and build out systems that contain fewer risks Services and microservices: understand the value of services for building complicated applications that need to operate at higher scale Scaling applications: assign services to specific teams, label the criticalness of each service, and devise failure scenarios and recovery plans Cloud services: understand the structure of cloud-based services, resource allocation, and service distribution

**ServiceNow Application Development** Jun 18 2020 Develop and extend efficient cloud-native applications with ServiceNow About This Book Build and customize your apps

and workflows to suit your organization's requirements Perform in-depth application development from designing forms to writing business rules, client-scripts, and workflows Comprehensive guide to the end-to-end implementation of designing and extending apps with ServiceNow Who This Book Is For If you are a ServiceNow administrator and developer and need to build and customize your service management solution (apps and workflows) with ServiceNow, then this book is for you. What You Will Learn Customize the ServiceNow dashboard to meet your business requirements Use Administration and Security Controls to add roles and ensure proper access Manage tables and columns using data dictionaries Learn how application scopes are defined within ServiceNow Configure different types of table to design your application Start using the different types of scripting options available in ServiceNow Design and create workflows for task tables Use debugging techniques available in ServiceNow to easily resolve script-related issues Run scripts at regular time intervals using the Scheduled Script Execution module In Detail ServiceNow provides service management for every department in the enterprise, including IT, Human Resources, Facilities, Field Service, and more. This book focuses on all the steps required to develop apps and workflows for any of your business requirements using ServiceNow. You will start with the first module, which covers the basics of ServiceNow and how applications are structured; how you can customize the dashboard as required; and also how to create users. After you get used to the dashboard, you will move on to the next module, Applications and Tables, where you will learn about working with different tables and how you can create a scope other than the global scope for your application. The next module is Scripting and APIs, where you will learn Scripting in ServiceNow and use powerful APIs to develop applications. The final module, Administration Essentials, covers debugging, advanced database features, and scheduled script creation. By the end of the book you will have mastered creating organized and customer-friendly applications Style and approach A step-by-step tutorial to designing applications and workflows with ServiceNow Apache Karaf Cookbook Apr 16 2020 This book is intended for developers who have some familiarity with Apache Karaf and who want a quick reference for practical, proven tips on how to perform common tasks such as configuring Pax modules deployed in Apache Karaf, Extending HttpService with Apache Karaf. You should have working knowledge of Apache karaf, as the book provides a deeper understanding of the capabilities of Apache Karaf.

**Microsoft Exchange Server 2013 Inside Out Mailbox and High Availability** May 18 2020 With a focus on mailbox and high availability features, this book delivers the ultimate, in-depth reference to IT professionals planning and managing an Exchange Server 2013 deployment. Guided by Tony Redmond, a Microsoft MVP and award-winning author, you will: Understand major changes to Exchange Server architecture Get inside insights for planning your upgrade or deployment Examine the new web-based Exchange admin center (EAC) Take a deep dive into configuring mailboxes, distribution groups, and contacts; planning and managing the Managed Store; database availability groups; mailbox replication service; compliance, data leakage, and data loss prevention; site mailboxes; modern public folders

Pro SQL Server Always On Availability Groups Oct 03 2021 This book is your field

guide to planning, deploying, managing, and troubleshooting Always On Availability Groups. The Always On Availability Groups feature is an enterprise-level solution for high availability and disaster recovery. Always On provides a rich set of options that reduce overhead and resource usage and enable database administrators to implement and manage high availability and disaster recovery solutions far more easily than by implementing the log-shipping and database-mirroring solutions of the past. Pro SQL Server Always On Availability Groups shows how to reduce downtime, maximize application availability, and provide data protection. The goal is always to have your SQL Server databases up and running whenever you need them, rain or shine, disaster or otherwise. With a focus on real-world experiences and war stories, authors Uttam Parui and Vivek Sanil offer you: Tips, tricks, and best practices for architecting and deploying availability groups. The confidence required to manage and monitor availability groups. Techniques to troubleshoot common issues that you may face during and after deploying availability groups in a mission-critical production environment. What You Will Learn Grasp important concepts underlying high-availability and disaster recovery. Plan and deploy Always On Availability Groups in your corporate environment. Manage Availability Groups to ensure constant readiness and high throughput. Monitor Availability Group performance and troubleshoot problems fast. Build on the cloud and make Windows Azure part of your availability solution. Employ proven techniques and best practices as tested and shared by the authors. Audience Pro SQL Server Always On Availability Groups is aimed at SQL Server architects, database administrators, and IT professionals who are tasked with architecting and deploying a high-availability and disaster recovery solution involving Microsoft SQL Server. This book is also for SQL Server support staff who will be managing and supporting existing availability group deployments.

**High Performance MySQL** Oct 11 2019 As users come to depend on MySQL, they find that they have to deal with issues of reliability, scalability, and performance--issues that are not well documented but are critical to a smoothly functioning site. This book is an insider's guide to these little understood topics. Author Jeremy Zawodny has managed large numbers of MySQL servers for mission-critical work at Yahoo!, maintained years of contacts with the MySQL AB team, and presents regularly at conferences. Jeremy and Derek have spent months experimenting, interviewing major users of MySQL, talking to MySQL AB, benchmarking, and writing some of their own tools in order to produce the information in this book. In High Performance MySQL you will learn about MySQL indexing and optimization in depth so you can make better use of these key features. You will learn practical replication, backup, and load-balancing strategies with information that goes beyond available tools to discuss their effects in real-life environments. And you'll learn the supporting techniques you need to carry out these tasks, including advanced configuration, benchmarking, and investigating logs. Topics include: A review of configuration and setup options Storage engines and table types Benchmarking Indexes Query Optimization Application Design Server Performance Replication Load-balancing Backup and Recovery Security

High Availability and Scalability of Mainframe Environments Using System Z and Z/OS as Example Feb 24 2021 Mainframe computers are the backbone of industrial and

commercial computing, hosting the most relevant and critical data of businesses. One of the most important mainframe environments is IBM System z with the operating system z/OS. This book introduces mainframe technology of System z and z/OS with respect to high availability and scalability. It highlights their presence on different levels within the hardware and software stack to satisfy the needs for large IT organizations.

**Designing High Availability Systems** Dec 25 2020 A practical, step-by-step guide to designing world-class, high availability systems using both classical and DFSS reliability techniques Whether designing telecom, aerospace, automotive, medical, financial, or public safety systems, every engineer aims for the utmost reliability and availability in the systems he, or she, designs. But between the dream of world-class performance and reality falls the shadow of complexities that can bedevil even the most rigorous design process. While there are an array of robust predictive engineering tools, there has been no single-source guide to understanding and using them . . . until now. Offering a case-based approach to designing, predicting, and deploying world-class high-availability systems from the ground up, this book brings together the best classical and DFSS reliability techniques. Although it focuses on technical aspects, this guide considers the business and market constraints that require that systems be designed right the first time. Written in plain English and following a step-by-step "cookbook" format, *Designing High Availability Systems*: Shows how to integrate an array of design/analysis tools, including Six Sigma, Failure Analysis, and Reliability Analysis Features many real-life examples and case studies describing predictive design methods, tradeoffs, risk priorities, "what-if" scenarios, and more Delivers numerous high-impact takeaways that you can apply to your current projects immediately Provides access to MATLAB programs for simulating problem sets presented, along with PowerPoint slides to assist in outlining the problem-solving process *Designing High Availability Systems* is an indispensable working resource for system engineers, software/hardware architects, and project teams working in all industries.

*MySQL High Availability* Jan 18 2023 Server bottlenecks and failures are a fact of life in any database deployment, but they don't have to bring everything to a halt. This practical book explains replication, cluster, and monitoring features that can help protect your MySQL system from outages, whether it's running on hardware, virtual machines, or in the cloud. Written by engineers who designed many of the tools covered, this book reveals undocumented or hard-to-find aspects of MySQL reliability and high availability—knowledge that's essential for any organization using this database system. This second edition describes extensive changes to MySQL tools. Versions up to 5.5 are covered, along with several 5.6 features. Learn replication fundamentals, including use of the binary log and MySQL Replicant Library Handle failing components through redundancy Scale out to manage read-load increases, and use data sharding to handle large databases and write-load increases Store and replicate data on individual nodes with MySQL Cluster Monitor database activity and performance, and major operating system parameters Keep track of masters and slaves, and deal with failures and restarts, corruption, and other incidents Examine tools including MySQL Enterprise Monitor, MySQL Utilities, and GTIDs

**Pro SharePoint Disaster Recovery and High Availability** Aug 01 2021 Few IT

professionals take the time to learn what needs to be known to do disaster recovery well. Most labor under the pretense that good administration equals close to five-nines uptime. Most technical people do not see the value of planning for disasters until the unexpected has already happened, and the effects of a disaster involving a SharePoint farm—which today houses business information, line-of-business applications, sensitive information, extranets, and other highly important assets—can be staggering. Pro SharePoint Disaster Recovery and High Availability, Second Edition will take you through a step-by-step process to show how to build an awareness and reaction plan for the inevitable. With a focus on real-world experiences and war stories, author Stephen Cummins weaves an expert tale of woe response and offers you: Ways to see the warning signs of disaster, and ways to avoid it Ways to respond to a disaster while it is happening Perhaps most importantly, how to develop a plan to deal with disaster when it inevitably does happen

**JUNOS High Availability** Mar 08 2022 Whether your network is a complex carrier or just a few machines supporting a small enterprise, JUNOS High Availability will help you build reliable and resilient networks that include Juniper Networks devices. With this book's valuable advice on software upgrades, scalability, remote network monitoring and management, high-availability protocols such as VRRP, and more, you'll have your network uptime at the five, six, or even seven nines -- or 99.99999% of the time. Rather than focus on "greenfield" designs, the authors explain how to intelligently modify multi-vendor networks. You'll learn to adapt new devices to existing protocols and platforms, and deploy continuous systems even when reporting scheduled downtime. JUNOS High Availability will help you save time and money. Manage network equipment with Best Common Practices Enhance scalability by adjusting network designs and protocols Combine the IGP and BGP networks of two merging companies Perform network audits Identify JUNOScripting techniques to maintain high availability Secure network equipment against breaches, and contain DoS attacks Automate network configuration through specific strategies and tools This book is a core part of the Juniper Networks Technical Library™.

*Teach Yourself Linux Virtualization and High Availability* Mar 16 2020 High availability server virtualization currently powers the vast majority of public-facing compute deployments and Linux lies at the heart of nearly all of them. If you aren't already engaged in a virtualized project that touches some kind of Linux technology, you probably will be soon. What are you doing to build your skills to meet the future? The Linux Professional Institute's LPIC-3 304 certification expectations are an excellent, vendor neutral introduction to Linux server virtualization and cluster management. Even if you don't have plans to take the exam and earn the certification itself right now, using the 304 as a curriculum guide is a smart move. And, one way or another, this book is a great primary resource.

**High Performance Drupal** Feb 13 2020 How can you help your Drupal website continue to perform at the highest level as it grows to meet demand? This comprehensive guide provides best practices, examples, and in-depth explanations for solving several performance and scalability issues. You'll learn how to apply coding and infrastructure techniques to Drupal internals, application performance, databases, web servers, and performance analysis. Covering Drupal versions 7 and 8, this book is the ideal reference



for everything from site deployment to implementing specific technologies such as Varnish, memcache, or Solr. If you have a basic understanding of Drupal and the Linux-Apache-MySQL-PHP (LAMP) stack, you're ready to get started. Establish a performance baseline and define goals for improvement Optimize your website's code and front-end performance Get best and worst practices for customizing Drupal core functionality Apply infrastructure design techniques to launch or expand a site Use tools to configure, monitor, and optimize MySQL performance Employ alternative storage and backend search options as your site grows Tune your web servers through httpd and PHP configuration Monitor services and perform load tests to catch problems before they become critical

**Oracle High Availability, Disaster Recovery, and Cloud Services** Aug 21 2020 Work with Oracle database's high-availability and disaster-management technologies. This book covers all the Oracle high-availability technologies in one place and also discusses how you configure them in engineered systems and cloud services. You will see that when you say your database is healthy, it is not limited to whether the database is performing well on day-to-day operations; rather it should also be robust and free from disasters. As a result, your database will be capable of handling unforeseen incidents and recovering from disaster with very minimal or zero downtime. Oracle High Availability, Disaster Recovery, and Cloud Services explores all the high-availability features of Oracle database, how to configure them, and best practices. After you have read this book you will have mastered database high-availability concepts such as RAC, Data Guard, OEM 13c, and engineered systems (Oracle Exadata x6/x7 and Oracle Database Appliance). What You Will Learn Master the best practices and features of Exadata and ODA Implement and monitor high availability with OEM 13c Clone databases using various methods in Oracle 12c R2 Work with the Oracle sharding features of Oracle 12c R2 Who This Book Is For Oracle database administrators

High Availability and Disaster Recovery Nov 16 2022 Companies and institutions depend more than ever on the availability of their Information Technology, and most mission critical business processes are IT-based. Business Continuity is the ability to do business under any circumstances and is an essential requirement faced by modern companies. Both concepts - High Availability and Disaster Recovery - are realized by redundant systems. This book presents requirements, concepts, and realizations of redundant systems on all abstraction levels, and all given examples refer to UNIX and Linux Systems.

**Juniper SRX Series** Nov 04 2021 This complete field guide, authorized by Juniper Networks, is the perfect hands-on reference for deploying, configuring, and operating Juniper's SRX Series networking device. Authors Brad Woodberg and Rob Cameron provide field-tested best practices for getting the most out of SRX deployments, based on their extensive field experience. While their earlier book, Junos Security, covered the SRX platform, this book focuses on the SRX Series devices themselves. You'll learn how to use SRX gateways to address an array of network requirements—including IP routing, intrusion detection, attack mitigation, unified threat management, and WAN acceleration. Along with case studies and troubleshooting tips, each chapter provides study questions and lots of useful illustrations. Explore SRX components, platforms, and various

deployment scenarios Learn best practices for configuring SRX's core networking features Leverage SRX system services to attain the best operational state Deploy SRX in transparent mode to act as a Layer 2 bridge Configure, troubleshoot, and deploy SRX in a highly available manner Design and configure an effective security policy in your network Implement and configure network address translation (NAT) types Provide security against deep threats with AppSecure, intrusion protection services, and unified threat management tools

*Cassandra High Availability* Jul 12 2022 If you are a developer or DevOps engineer who understands the basics of Cassandra and are ready to take your knowledge to the next level, then this book is for you. An understanding of the essentials of Cassandra is needed.

**Architecting for Scale** May 30 2021 Every day, companies struggle to scale critical applications. As traffic volume and data demands increase, these applications become more complicated and brittle, exposing risks and compromising availability. With the popularity of software as a service, scaling has never been more important. Updated with an expanded focus on modern architecture paradigms such as microservices and cloud computing, this practical guide provides techniques for building systems that can handle huge quantities of traffic, data, and demand—without affecting the quality your customers expect. Architects, managers, and directors in engineering and operations organizations will learn how to build applications at scale that run more smoothly and reliably to meet the needs of customers. Learn how scaling affects the availability of your services, why that matters, and how to improve it Dive into a modern service-based application architecture that ensures high availability and reduces the effects of service failures Explore the Single Team Owned Service Architecture paradigm (STOSA)—a model for scaling your development organization in tandem with your application Understand, measure, and mitigate risk in your systems Use the cloud to build highly scalable applications

**Microsoft SQL Server 2000 High Availability** Feb 07 2022 Get the definitive reference to maximizing the availability of your critical, SQL Server 2000-based business systems. Developed by Microsoft field consultants who create and support high availability solutions every day, this practical guide explains how to plan, design, implement, and support highly available SQL Server 2000-based systems for both 32-bit and 64-bit environments. You get in-depth, end-to-end guidance—from evaluating availability requirements and capacity planning to implementing failover clustering, log shipping, replication, system backup and restore, system recovery, and other technologies and processes. Real-world scenarios illustrate the book's best practices in action. You also get a CD featuring evaluation software, planning worksheets, scripts, and other on-the-job resources. Discover how to: Evaluate technical, staffing, and budget requirements Plan and configure disk subsystems for performance and availability Configure Microsoft Windows 2000 Server or Microsoft Windows Server 2003 to optimize SQL Server 2000 availability Plan, implement, and maintain a base server cluster and Network Load Balancing Plan and configure failover clustering Use built-in log shipping features or code your own Create comprehensive backup and restore plans Choose replication methods for high availability—and expose replication architecture Develop prevention

and disaster recovery strategies; sample runbook and scripts included on CD Minimize downtime for server upgrades and service pack installations Baseline performance and use availability-monitoring tools and techniques CDs feature: 120-day evaluation version of SQL Server 2000 Enterprise Edition Scripts for custom log shipping, backup and restoration, system monitoring, and other administrative tasks Sample disaster-recovery runbook Worksheets, checklists, and templates Fully searchable eBook For customers who purchase an ebook version of this title, instructions for downloading the CD files can be found in the ebook.

*Professional SQL Server High Availability and Disaster Recovery* Jun 30 2021 Leverage powerful features of the SQL Server and watch your infrastructure transform into a high-performing, reliable network of systems. Key Features Explore more than 20 real-world use cases to understand SQL Server features Get to grips with the SQL Server Always On technology Learn how to choose HA and DR topologies for your system Book Description Professional SQL Server High Availability and Disaster Recovery explains the high availability and disaster recovery technologies available in SQL Server: Replication, AlwaysOn, and Log Shipping. You'll learn what they are, how to monitor them, and how to troubleshoot any related problems. You will be introduced to the availability groups of AlwaysOn and learn how to configure them to extend your database mirroring. Through this book, you will be able to explore the technical implementations of high availability and disaster recovery technologies that you can use when you create a highly available infrastructure, including hybrid topologies. Note that this course does not cover SQL Server Failover Cluster Installation with shared storage. By the end of the book, you'll be equipped with all that you need to know to develop robust and high performance infrastructure. What you will learn Configure and troubleshoot Replication, AlwaysOn, and Log Shipping Study the best practices to implement HA and DR solutions Design HA and DR topologies for the SQL Server and study how to choose a topology for your environment Use T-SQL to configure replication, AlwaysOn, and log shipping Migrate from On-Premise SQL Server to Azure SQL Database Manage and maintain AlwaysOn availability groups for extended database mirroring Who this book is for Professional SQL Server High Availability and Disaster Recovery is for you if you are a database administrator or database developer who wants to improve the performance of your production environment. Prior experience of working with SQL Server will help you get the most out of this book.

**Pro Linux High Availability Clustering** Dec 17 2022 Pro Linux High Availability Clustering teaches you how to implement this fundamental Linux add-on into your business. Linux High Availability Clustering is needed to ensure the availability of mission critical resources. The technique is applied more and more in corporate datacenters around the world. While lots of documentation about the subject is available on the internet, it isn't always easy to build a real solution based on that scattered information, which is often oriented towards specific tasks only. Pro Linux High Availability Clustering explains essential high-availability clustering components on all Linux platforms, giving you the insight to build solutions for any specific case needed. In this book four common cases will be explained: Configuring Apache for high availability Creating an Open Source SAN based on DRBD, iSCSI and HA clustering Setting up a

load-balanced web server cluster with a back-end, highly-available database Setting up a KVM virtualization platform with high-availability protection for a virtual machine. With the knowledge you'll gain from these real-world applications, you'll be able to efficiently apply Linux HA to your work situation with confidence. Author Sander Van Vugt teaches Linux high-availability clustering on training courses, uses it in his everyday work, and now brings this knowledge to you in one place, with clear examples and cases. Make the best start with HA clustering with Pro Linux High Availability Clustering at your side.

Blueprints for High Availability Oct 15 2022 Expert techniques for designing your system to achieve maximum availability and predictable downtime With your company's reputation and profits at stake, downtime on your 24/7 web site is not an option, nor is poor application performance. Now in its second edition, this authoritative book provides you with the design blueprints to maximize your system availability. Striking a balance between costs and benefits, the authors show you all of the elements of your computer system that can fail-as well as ways to assess their reliability and attain resiliency and high availability for each one. A unique feature is "Tales from the Field," a collection of true-to-life experiences that will help you avoid mistakes and deploy your system with confidence. Learn how to design your system to limit the impact of such problems as computer viruses, natural disasters, or the corruption of critical files and discover how to:

- \* Implement effective backup-and-restore and tape management strategies
- \* Arrange disks and disk arrays to avoid downtime caused by inevitable failures
- \* Utilize technologies such as Storage Area Networks (SANs), Network Attached Storage (NAS), Virtualization, and clustering
- \* Achieve effective application recovery after any part of the system has failed
- \* Replicate critical data to remote systems across a network

*NGINX Cookbook* Aug 13 2022 NGINX is one of the most widely used web servers available today, in part because of its capabilities as a load balancer and reverse proxy server for HTTP and other network protocols. This cookbook provides easy-to-follow examples to real-world problems in application delivery. The practical recipes will help you set up and use either the open source or commercial offering to solve problems in various use cases. For professionals who understand modern web architectures, such as n-tier or microservice designs, and common web protocols including TCP and HTTP, these recipes provide proven solutions for security, software load balancing, and monitoring and maintaining NGINX's application delivery platform. You'll also explore advanced features of both NGINX and NGINX Plus, the free and licensed versions of this server. You'll find recipes for: High-performance load balancing with HTTP, TCP, and UDP Securing access through encrypted traffic, secure links, HTTP authentication subrequests, and more Deploying NGINX to Google Cloud, AWS, and Azure cloud computing services Setting up and configuring NGINX Controller Installing and configuring the NGINX Plus App Protect module Enabling WAF through Controller ADC

**CentOS High Availability** Dec 05 2021 This book is targeted at system engineers and system administrators who want to upgrade their knowledge and skills in high availability and want to learn practically how to achieve high availability with CentOS Linux. You are expected to have good CentOS Linux knowledge and basic networking experience.

Blueprints for High Availability May 10 2022 Expert techniques for designing your system to achieve maximum availability and predictable downtime With your company's

reputation and profits at stake, downtime on your 24/7 web site is not an option, nor is poor application performance. Now in its second edition, this authoritative book provides you with the design blueprints to maximize your system availability. Striking a balance between costs and benefits, the authors show you all of the elements of your computer system that can fail-as well as ways to assess their reliability and attain resiliency and high availability for each one. A unique feature is "Tales from the Field," a collection of true-to-life experiences that will help you avoid mistakes and deploy your system with confidence. Learn how to design your system to limit the impact of such problems as computer viruses, natural disasters, or the corruption of critical files and discover how to:

- \* Implement effective backup-and-restore and tape management strategies
- \* Arrange disks and disk arrays to avoid downtime caused by inevitable failures
- \* Utilize technologies such as Storage Area Networks (SANs), Network Attached Storage (NAS), Virtualization, and clustering
- \* Achieve effective application recovery after any part of the system has failed
- \* Replicate critical data to remote systems across a network

High Availability and Disaster Recovery Options for DB2 for Linux, UNIX, and Windows Sep 02 2021 As organizations strive to do more with less, IBM® DB2® for Linux, UNIX, and Windows provides various built-in high availability features. DB2 further provides high availability solutions by using enterprise system resources with broad support for clustering software, such as IBM PowerHA® SystemMirror®, IBM Tivoli® System Automation for Multiplatforms (Tivoli SA MP), and Microsoft Windows Cluster Server. This IBM Redbooks® publication describes the DB2 high availability functions and features, focusing on High Availability Disaster Recovery (HADR) in the OLTP environment. The book provides a detailed description of HADR, including setup, configuration, administration, monitoring, and preferred practices. This book explains how to configure Cluster software PowerHA, Tivoli SA MP, and MSCS with DB2 and show how to use these products to automate HADR takeover. DB2 also provides unprecedented enterprise-class disaster recovery capability. This book covers single system view backup, backup and restore with snapshot backup, and the db2recovery command, in detail. This book is intended for database administrators and information management professionals who want to design, implement, and support a highly available DB2 system.

**End-to-End High Availability Solution for System z from a Linux Perspective** Jan 06 2022 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. This IBM Redbooks publication starts with an explanation of high availability (HA) fundamentals such as HA concepts and terminology. It continues with a discussion of why a business needs to consider an HA solution and then explains how to determine your business single points of failure. We outline the components of a high availability solution and describe these components. Then we provide some architectural scenarios and demonstrate how to plan and decide an implementation of an end-to-end HA solution, from Linux on System z database scenarios to z/OS, and include storage, network, z/VM, Linux, and middleware. This implementation includes the IBM Tivoli System Automation for Multiplatforms (TSA MP), which monitors and automates applications distributed across Linux, AIX®, and z/OS® operating systems, as well as a GDPS based solution. It includes the planning

for an end-to-end scenario, considering Linux on System z, z/VM, and z/OS operating environments, and the middleware used. The TSA MP implements HA for infrastructure, network, operating systems, and applications across multiple platforms and is compared to a Linux HA implementation based on open source Linux-HA, which is Linux only.

*High Availability* Feb 19 2023 A best practices guide to the people and process issues associated with maximizing application availability. Focus is on how enterprises can design systems that are easier to maintain.

JBoss EAP6 High Availability Jan 14 2020 An easy-to-follow guide full of hands-on examples of real-world administration tasks. JBoss EAP6 High Availability is ideal for those who want to learn how to use JBoss EAP6 to set up a cluster. Basic knowledge of Linux/Unix is required.

**PostgreSQL High Availability Cookbook - Second Edition** Nov 23 2020 Over 100 recipes to design and implement a highly available server with the advanced features of PostgreSQL 9.4, 9.5 and 9.6

About This Book\* Create a PostgreSQL cluster that stays online even when disaster strikes\* Avoid costly downtime and data loss that can ruin your business\* Updated to include the newest features introduced in PostgreSQL 9.6 with hands-on industry-driven recipes

Who This Book Is For If you are a PostgreSQL DBA working on Linux systems who want a database that never gives up, this book is for you. If you've ever experienced a database outage, restored from a backup, spent hours trying to repair a malfunctioning cluster, or simply want to guarantee system stability, this book is definitely for you.

What you will learn\* Protect your data with PostgreSQL replication and management tools such as Slony, Bucardo, pglogical, and WAL-E\* Hardware planning to help your database run efficiently\* Prepare for catastrophes and prevent them before they happen\* Reduce database resource contention with connection pooling using pgpool and PgBouncer\* Automate monitoring and alerts to visualize cluster activity using Nagios and collected\* Construct a robust software stack that can detect and fix outages\* Learn simple PostgreSQL High Availability with Patroni, or dive into the full power of Pacemaker.

In Detail Databases are nothing without the data they store. In the event of a failure - catastrophic or otherwise - immediate recovery is essential. By carefully combining multiple servers, it's even possible to hide the fact a failure occurred at all. From hardware selection to software stacks and horizontal scalability, this book will help you build a versatile PostgreSQL cluster that will survive crashes, resist data corruption, and grow smoothly with customer demand. It all begins with hardware selection for the skeleton of an efficient PostgreSQL database cluster. Then it's on to preventing downtime as well as troubleshooting some real life problems that administrators commonly face. Next, we add database monitoring to the stack, using collectd, Nagios, and Graphite. And no stack is complete without replication using multiple internal and external tools, including the newly released pglogical extension. Pacemaker or Raft consensus tools are the final piece to grant the cluster the ability to heal itself. We even round off by tackling the complex problem of data scalability. This book exploits many new features introduced in PostgreSQL 9.6 to make the database more efficient and adaptive, and most importantly, keep it running.

**High Availability MySQL Cookbook** Jan 26 2021 Over 60 simple but incredibly effective recipes focusing on different methods of achieving high availability for MySQL

database.

**Big Data Analytics with Hadoop 3** Nov 11 2019 Explore big data concepts, platforms, analytics, and their applications using the power of Hadoop 3 Key Features Learn Hadoop 3 to build effective big data analytics solutions on-premise and on cloud Integrate Hadoop with other big data tools such as R, Python, Apache Spark, and Apache Flink Exploit big data using Hadoop 3 with real-world examples Book Description Apache Hadoop is the most popular platform for big data processing, and can be combined with a host of other big data tools to build powerful analytics solutions. Big Data Analytics with Hadoop 3 shows you how to do just that, by providing insights into the software as well as its benefits with the help of practical examples. Once you have taken a tour of Hadoop 3's latest features, you will get an overview of HDFS, MapReduce, and YARN, and how they enable faster, more efficient big data processing. You will then move on to learning how to integrate Hadoop with the open source tools, such as Python and R, to analyze and visualize data and perform statistical computing on big data. As you get acquainted with all this, you will explore how to use Hadoop 3 with Apache Spark and Apache Flink for real-time data analytics and stream processing. In addition to this, you will understand how to use Hadoop to build analytics solutions on the cloud and an end-to-end pipeline to perform big data analysis using practical use cases. By the end of this book, you will be well-versed with the analytical capabilities of the Hadoop ecosystem. You will be able to build powerful solutions to perform big data analytics and get insight effortlessly. What you will learn Explore the new features of Hadoop 3 along with HDFS, YARN, and MapReduce Get well-versed with the analytical capabilities of Hadoop ecosystem using practical examples Integrate Hadoop with R and Python for more efficient big data processing Learn to use Hadoop with Apache Spark and Apache Flink for real-time data analytics Set up a Hadoop cluster on AWS cloud Perform big data analytics on AWS using Elastic Map Reduce Who this book is for Big Data Analytics with Hadoop 3 is for you if you are looking to build high-performance analytics solutions for your enterprise or business using Hadoop 3's powerful features, or you're new to big data analytics. A basic understanding of the Java programming language is required.

*IBM High Availability Solution for IBM FileNet P8 Systems* Mar 28 2021

**SQL Server 2016 High Availability Unleashed (includes Content Update Program)**

Apr 28 2021 Book + Content Update Program SQL Server 2016 High Availability Unleashed provides start-to-finish coverage of SQL Server's powerful high availability (HA) solutions for your traditional on-premise databases, cloud-based databases (Azure or AWS), hybrid databases (on-premise coupled with the cloud), and your emerging Big Data solutions. This complete guide introduces an easy-to-follow, formal HA methodology that has been refined over the past several years and helps you identify the right HA solution for your needs. There is also additional coverage of both disaster recovery and business continuity architectures and considerations. You are provided with step-by-step guides, examples, and sample code to help you set up, manage, and administer these highly available solutions. All examples are based on existing production deployments at major Fortune 500 companies around the globe. This book is for all intermediate-to-advanced SQL Server and Big Data professionals, but is also

organized so that the first few chapters are great foundation reading for CIOs, CTOs, and even some tech-savvy CFOs. Learn a formal, high availability methodology for understanding and selecting the right HA solution for your needs Deep dive into Microsoft Cluster Services Use selective data replication topologies Explore thorough details on AlwaysOn and availability groups Learn about HA options with log shipping and database mirroring/ snapshots Get details on Microsoft Azure for Big Data and Azure SQL Explore business continuity and disaster recovery Learn about on-premise, cloud, and hybrid deployments Provide all types of database needs, including online transaction processing, data warehouse and business intelligence, and Big Data Explore the future of HA and disaster recovery In addition, this book is part of InformIT's exciting Content Update Program, which provides content updates for major technology improvements! As significant updates are made to SQL Server, sections of this book will be updated or new sections will be added to match the updates to the technologies. As updates become available, they will be delivered to you via a free Web Edition of this book, which can be accessed with any Internet connection. To learn more, visit [informit.com/cup](http://informit.com/cup). How to access the Web Edition: Follow the instructions inside to learn how to register your book to access the FREE Web Edition.

**High Availability Network Fundamentals** Dec 13 2019 A practical guide to modeling and designing reliable networks Provides a detailed introduction to modeling availability necessary for network design Helps network designers understand the theoretical availability of their topologies Explains the factors that limit availability to minimize the number of network failures Provides all the information necessary to do basic availability modeling/budgeting High Availability Network Fundamentalsdiscusses the need for and the mathematics of availability, then moves on to cover the issues affecting availability, including hardware, software, design strategies, human error, and environmental considerations. After setting up the range of common problems, it then delves into the details of how to design networks for fault tolerance and provides sample calculations for specific systems. Also included is a complete, end-to-end example showing availability calculations for a sample network.

- [High Availability](#)
- [MySQL High Availability](#)
- [Pro Linux High Availability Clustering](#)
- [High Availability And Disaster Recovery](#)
- [Blueprints For High Availability](#)
- [MySQL High Availability](#)
- [NGINX Cookbook](#)
- [Cassandra High Availability](#)
- [Clusters For High Availability](#)
- [Blueprints For High Availability](#)
- [Understanding And Using Q Replication For High Availability Solutions On The IBM Z OS Platform](#)
- [JUNOS High Availability](#)



- [Microsoft SQL Server 2000 High Availability](#)
- [End to End High Availability Solution For System Z From A Linux Perspective](#)
- [CentOS High Availability](#)
- [Juniper SRX Series](#)
- [Pro SQL Server Always On Availability Groups](#)
- [High Availability And Disaster Recovery Options For DB2 For Linux UNIX And Windows](#)
- [Pro SharePoint Disaster Recovery And High Availability](#)
- [Professional SQL Server High Availability And Disaster Recovery](#)
- [Architecting For Scale](#)
- [SQL Server 2016 High Availability Unleashed Includes Content Update Program](#)
- [IBM High Availability Solution For IBM FileNet P8 Systems](#)
- [High Availability And Scalability Of Mainframe Environments Using System Z And Z OS As Example](#)
- [High Availability MySQL Cookbook](#)
- [Designing High Availability Systems](#)
- [PostgreSQL High Availability Cookbook Second Edition](#)
- [Linux In Action](#)
- [Architecting For Scale](#)
- [Oracle High Availability Disaster Recovery And Cloud Services](#)
- [PostgreSQL 12 High Availability Cookbook](#)
- [ServiceNow Application Development](#)
- [Microsoft Exchange Server 2013 Inside Out Mailbox And High Availability](#)
- [Apache Karaf Cookbook](#)
- [Teach Yourself Linux Virtualization And High Availability](#)
- [High Performance Drupal](#)
- [JBoss EAP6 High Availability](#)
- [High Availability Network Fundamentals](#)
- [Big Data Analytics With Hadoop 3](#)
- [High Performance MySQL](#)