

Monitoring With Nagios And Check Mk

The future for Nagios in the enterprise is certainly bright! Nagios 3 Enterprise Network Monitoring can help you harness the full power of Nagios in your organization. Nagios 3 contains many significant new features and updates, and this book details them all for you. Once up and running, you'll see how a number of useful add-ons and enhancements for Nagios can extend the functionality of Nagios throughout your organization. And, if you want to learn how to write your own plugins...this is the book for you! In these pages you'll find a cookbook-style chapter full of useful plugins that monitor a variety of devices, from HTTP-based applications to CPU utilization to LDAP servers and more. Complete Case Study Demonstrates how to Deploy Nagios Globally in an Enterprise Network Monitor Third Party Hardware Devices with Nagios

The Fully Updated Guide to Enterprise Network Monitoring with Today's Nagios Platform and Tools This is the definitive guide to building cost-effective, enterprise-strength monitoring infrastructures with the latest commercial and open source versions of Nagios. World-renowned monitoring expert David Josephsen covers the entire monitoring software stack, treating Nagios as a specification language and foundation for building well designed monitoring systems that can scale to serve any organization. Drawing on his unsurpassed experience, Josephsen demonstrates best practices throughout and also reveals common mistakes, their consequences, and how to avoid them. He provides all the technical depth you need to configure and run Nagios successfully, including a practical and thorough discussion of writing your own custom modules with the C-based Nagios Event-Broker API. Extensively updated throughout, this edition adds an entirely new chapter on scaling Nagios for large, complex networks that rely heavily on virtualization and cloud services. Josephsen thoroughly introduces Nagios XI, the advanced new commercial version of Nagios and shows how to improve productivity with the latest third-party tools and plug-ins. Coverage includes: Learn how Nagios works, in depth Master focused, efficient techniques for configuring and deploying the latest versions of Nagios Solve real-world problems in monitoring Windows and UNIX systems, networking hardware, and environmental sensors Systematically scale and optimize Nagios for the largest enterprise environments Enhance your monitoring system with new tools including Check-MK, Op5 Merlin, and SFlow Integrate visualization via Ganglia, Graphite, and RRDTool Simplify and streamline all facets of system monitoring with Nagios XI Build powerful custom Nagios Event Broker (NEB) modules, step-by-step Learn about easy-to-understand code listings, fully updated for today's platforms No matter how complex your systems monitoring challenges are, this book will help you achieve the results you want—right from the start.

This book is useful for Hadoop administrators who need to learn how to monitor and diagnose their clusters. Also, the book will prove useful for new users of the technology, as the language used is simple and easy to grasp.

How well does your enterprise stand up against today's sophisticated security threats? In this book, security experts from Cisco Systems demonstrate how to detect damaging security incidents on your global network--first by teaching you which assets you need to monitor closely, and then by helping you develop targeted strategies and pragmatic techniques to protect them. Security

Monitoring is based on the authors' years of experience conducting incident response to keep Cisco's global network secure. It offers six steps to improve network monitoring. These steps will help you: Develop Policies: define rules, regulations, and monitoring criteria Know Your Network: build knowledge of your infrastructure with network telemetry Select Your Targets: define the subset of infrastructure to be monitored Choose Event Sources: identify event types needed to discover policy violations Feed and Tune: collect data, generate alerts, and tune systems using contextual information Maintain Dependable Event Sources: prevent critical gaps in collecting and monitoring events Security Monitoring illustrates these steps with detailed examples that will help you learn to select and deploy the best techniques for monitoring your own enterprise network.

The author focuses solely on how UNIX and Linux system administrators can use well-known tools to automate tasks, even across multiple systems.

"This reference presents a vital compendium of research detailing the latest case studies, architectures, frameworks, methodologies, and research on Grid and Cloud Computing"--

This book aims to provide a deep look into Italian actions taken in some fields of science and high performance computing (HPC), and the Italian effort to bridge the HPC gap with respect to Europe. The Italian PON ReCaS Project is written for graduate readers and professionals in the field of high performance computing. It presents and discusses innovative and important technological solutions, and describes interesting results in various fields of application. ReCaS stands for "Rete di Calcolo per SuperB e altre applicazioni" and is a computing network infrastructure in Southern Italy devoted to scientific and non-scientific applications within the vision of a common European infrastructure for computing, storage and network. The ReCaS project is part of the 2007–2013 European Union strategy, and was funded by the Italian Ministry of Research and Education (MIUR) for the development and enhancement of a distributed computing infrastructure of the Grid/Cloud type over the four EU 'Convergence' regions in Southern Italy: Campania, Puglia and Sicily and Calabria. The network will be open and accessible to all researchers, public and private, and will be characterized by unprecedented computing power and storage capacity. Posted in the European Grid Infrastructure EGI, ReCaS is also an opportunity to the countries of the Mediterranean area and extends the potential of the current network. As part of Packt's cookbook series, each recipe offers a practical, step-by-step solution to common problems found in HBase administration. This book is for HBase administrators, developers, and will even help Hadoop administrators. You are not required to have HBase experience, but are expected to have a basic understanding of Hadoop and MapReduce.

Provides information on how to use Pro Nagios 3.0 to monitor and report on servers, network devices, and applications.

Network monitoring can be a complex task to implement and maintain in your IT infrastructure. Nagios, an open-source host, service and network monitoring program can help you streamline your network monitoring tasks and reduce the cost of operation. With this shortcut guide, we'll go over how Nagios fits in the overall network monitoring puzzle. We'll also cover installation and basic usage. Finally, we'll show you how to extend Nagios with other tools to extend functionality.

A comprehensive configuration guide to monitor and maintain your network and systems

Master over 100 recipes to design and implement a highly available server with the advanced features of PostgreSQL
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. Style and approach This book contains practical recipes that will help the reader solve real world problems related to high availability in PostgreSQL. Every recipe is explained in detail, with relevant explanations, tips and tricks provided for quicker and easier understanding.

Cloud computing has revolutionized computer systems, providing greater dynamism and flexibility to a variety of operations. It can help businesses quickly and effectively adapt to market changes, and helps promote users' continual access to vital information across platforms and devices. Cloud Computing Advancements in Design, Implementation, and Technologies outlines advancements in the state-of-the-art, standards, and practices of cloud computing, in an effort

to identify emerging trends that will ultimately define the future of the cloud. A valuable reference for academics and practitioners alike, this title covers topics such as virtualization technology, utility computing, cloud application services (SaaS), grid computing, and services computing.

This IBM Redbooks publication describes how to implement an Open Platform for Database as a Service (DBaaS) on IBM Power Systems environment for Linux, and demonstrate the open source tools, optimization and best practices guidelines for it. Open Platform for DBaaS on Power Systems is an on-demand, secure, and scalable self-service database platform that automates provisioning and administration of databases to support new business applications and information insights. This publication addresses topics to help sellers, architects, brand specialists, distributors, resellers and anyone offering secure and scalable Open Platform for DBaaS on Power Systems solution with APIs that are consistent across heterogeneous open database types. An Open Platform for DBaaS on Power Systems solution has the capability to accelerate business success by providing an infrastructure, and tools leveraging Open Source and OpenStack software engineered to optimize hardware and software between workloads and resources so you have a responsive, and an adaptive environment. Moreover, this publication provides documentation to transfer the how-to-skills for cloud oriented operational management of Open Platform for DBaaS on Power Systems service and underlying infrastructure to the technical teams. Open Platform for DBaaS on Power Systems mission is to provide scalable and reliable cloud database as a service provisioning functionality for both relational and non-relational database engines, and to continue to improve its fully-featured and extensible open source framework. For example, Trove is a database as a service for OpenStack. It is designed to run entirely on OpenStack, with the goal of allowing users to quickly and easily utilize the features of a relational or non-relational database without the burden of handling complex administrative tasks. Cloud users and database administrators can provision and manage multiple database instances as needed. Initially, the service focuses on providing resource isolation at high performance while automating complex administrative tasks including deployment, configuration, patching, backups, restores, and monitoring. In the context of this publication, the monitoring tool implemented is Nagios Core which is an open source monitoring tool. Hence, when you see a reference of Nagios in this book, Nagios Core is the open source monitoring solution implemented. Also note that the implementation of Open Platform for DBaaS on IBM Power Systems is based on open source solutions. This book is targeted toward sellers, architects, brand specialists, distributors, resellers and anyone developing and implementing Open Platform for DBaaS on Power Systems solutions.

A hands-on guide to leveraging NoSQL databases NoSQL databases are an efficient and powerful tool for storing and manipulating vast quantities of data. Most NoSQL databases scale well as data grows. In addition, they are often

malleable and flexible enough to accommodate semi-structured and sparse data sets. This comprehensive hands-on guide presents fundamental concepts and practical solutions for getting you ready to use NoSQL databases. Expert author Shashank Tiwari begins with a helpful introduction on the subject of NoSQL, explains its characteristics and typical uses, and looks at where it fits in the application stack. Unique insights help you choose which NoSQL solutions are best for solving your specific data storage needs. Professional NoSQL: Demystifies the concepts that relate to NoSQL databases, including column-family oriented stores, key/value databases, and document databases. Delves into installing and configuring a number of NoSQL products and the Hadoop family of products. Explains ways of storing, accessing, and querying data in NoSQL databases through examples that use MongoDB, HBase, Cassandra, Redis, CouchDB, Google App Engine Datastore and more. Looks at architecture and internals. Provides guidelines for optimal usage, performance tuning, and scalable configurations. Presents a number of tools and utilities relating to NoSQL, distributed platforms, and scalable processing, including Hive, Pig, RRDtool, Nagios, and more.

Build intelligent software stacks with the Puppet configuration management suite About This Book Develop high-quality Puppet modules in an isolated development environment Manage complex development environments with industry-leading configuration management tool A comprehensive guide to get you started with Puppet development and deployment in virtual environments Who This Book Is For If you are new to configuration management and IT automation processes and are looking for better ways to manage system configuration changes at scale, this book is for you. Basic knowledge of Linux System Administration is a prerequisite. What You Will Learn Manage your system with Puppet instantly Develop Puppet in an isolated development environment Make your manifests reusable to avoid re-inventing the wheel Automate monitoring to improve the user experience through increased uptime Enable nodes to communicate with each other via Puppet Master Make environment configuration dynamic using stored configurations and PuppetDB Extend Puppet beyond the built-in functionalities Manage your environment through the Puppet Enterprise console In Detail Puppet is a cross-platform, open source configuration management utility, which runs on various Unix, Linux, and Windows Microsoft platforms. It allows you to automate all your IT configurations, giving you control of what you do to each node, and also when and how you do it. You'll be able to build and manage development, test, and production environments independently without requiring previous system administration experience. Learning Puppet is a step-by-step guide on how to get started with Puppet development and use Puppet modules as the building blocks to deploy production-ready application cluster in virtual environment. You will begin with the installation of development environment on the VirtualBox hypervisor and Puppet Learning VM that will be used as the platform for testing and development of Puppet modules. Next, you will learn how to manage virtual machines and snapshots effectively and

enhance the development experience with advanced VirtualBox features. Later the book will focus on Puppet module development in detail. You will be guided through the process of utilizing existing modules that are available in the public module repository, write your own modules and use them to deploy a real-world web application that includes features such as monitoring and load balancing. You will then learn to scale your environment and turn your static configuration into a dynamic one through stored configurations and PuppetDB. Finally, the book will provide you with practical advice on Puppet troubleshooting and managing your environment with the wealth of features provided by the Puppet Enterprise console. Style and approach A comprehensive introductory guide to help you manage your infrastructure with Puppet. All instructions and explanations are supported with screenshots and code examples to ensure you get an easy start with Puppet.

CentOS is just like Red Hat, but without the price tag and with the virtuous license. When belts have to be tightened, we want to read about an OS with all the features of a commercial Linux variety, but without the pain. The Definitive Guide to CentOS is the first definitive reference for CentOS and focuses on CentOS alone, the workhorse Linux distribution, that does the heavy lifting in small and medium-size enterprises without drawing too much attention to itself. Provides tutorial and hands-on learning but is also designed to be used as a reference Bases all examples on real-world tasks that readers are likely to perform Serves up hard-won examples and hints and tips from the author's experiences of CentOS in production

Pro Puppet, Second Edition, now updated for Puppet 3, is an in-depth guide to installing, using, and developing the popular configuration management tool Puppet. Puppet provides a way to automate everything from user management to server configuration. You'll learn how Puppet has changed in the latest version, how to use it on a variety of platforms, including Windows, how to work with Puppet modules, and how to use Hiera. Puppet is a must-have tool for system administrators, and Pro Puppet will teach you how to maximize its capabilities and customize it for your environment. Install and configure Puppet to immediately start automating tasks and create reporting solutions Learn insider tricks and techniques to better manage your infrastructure Become a Puppet expert!

The Critical Infrastructure Protection Survey recently released by Symantec found that 53% of interviewed IT security experts from international companies experienced at least ten cyber attacks in the last five years, and financial institutions were often subject to some of the most sophisticated and large-scale cyber attacks and frauds. The book by Baldoni and Chockler analyzes the structure of software infrastructures found in the financial domain, their vulnerabilities to cyber attacks and the existing protection mechanisms. It then shows the advantages of sharing information among financial players in order to detect and quickly react to cyber attacks. Various aspects associated with information sharing

are investigated from the organizational, cultural and legislative perspectives. The presentation is organized in two parts: Part I explores general issues associated with information sharing in the financial sector and is intended to set the stage for the vertical IT middleware solution proposed in Part II. Nonetheless, it is self-contained and details a survey of various types of critical infrastructure along with their vulnerability analysis, which has not yet appeared in a textbook-style publication elsewhere. Part II then presents the CoMiFin middleware for collaborative protection of the financial infrastructure. The material is presented in an accessible style and does not require specific prerequisites. It appeals to both researchers in the areas of security, distributed systems, and event processing working on new protection mechanisms, and practitioners looking for a state-of-the-art middleware technology to enhance the security of their critical infrastructures in e.g. banking, military, and other highly sensitive applications. The latter group will especially appreciate the concrete usage scenarios included.

Pro Python System Administration, Second Edition explains and shows how to apply Python scripting in practice. It will show you how to approach and resolve real-world issues that most system administrators will come across in their careers. This book has been updated using Python 2.7 and Python 3 where appropriate. It also uses various new and relevant open source projects and tools that should now be used in practice. In this updated edition, you will find several projects in the categories of network administration, web server administration, and monitoring and database management. In each project, the author will define the problem, design the solution, and go through the more interesting implementation steps. Each project is accompanied by the source code of a fully working prototype, which you'll be able to use immediately or adapt to your requirements and environment. This book is primarily aimed at experienced system administrators whose day-to-day tasks involve looking after and managing small-to-medium-sized server estates. It will also be beneficial for system administrators who want to learn more about automation and want to apply their Python knowledge to solve various system administration problems. Python developers will also benefit from reading this book, especially if they are involved in developing automation and management tools.

Learn the right cutting-edge skills and knowledge to leverage Spark Streaming to implement a wide array of real-time, streaming applications. This book walks you through end-to-end real-time application development using real-world applications, data, and code. Taking an application-first approach, each chapter introduces use cases from a specific industry and uses publicly available datasets from that domain to unravel the intricacies of production-grade design and implementation. The domains covered in Pro Spark Streaming include social media, the sharing economy, finance, online advertising, telecommunication, and IoT. In the last few years, Spark has become synonymous with big data processing. DStreams enhance the underlying Spark processing engine to support streaming analysis with a novel micro-

batch processing model. Pro Spark Streaming by Zubair Nabi will enable you to become a specialist of latency sensitive applications by leveraging the key features of DStreams, micro-batch processing, and functional programming. To this end, the book includes ready-to-deploy examples and actual code. Pro Spark Streaming will act as the bible of Spark Streaming. What You'll Learn Discover Spark Streaming application development and best practices Work with the low-level details of discretized streams Optimize production-grade deployments of Spark Streaming via configuration recipes and instrumentation using Graphite, collectd, and Nagios Ingest data from disparate sources including MQTT, Flume, Kafka, Twitter, and a custom HTTP receiver Integrate and couple with HBase, Cassandra, and Redis Take advantage of design patterns for side-effects and maintaining state across the Spark Streaming micro-batch model Implement real-time and scalable ETL using data frames, SparkSQL, Hive, and SparkR Use streaming machine learning, predictive analytics, and recommendations Mesh batch processing with stream processing via the Lambda architecture Who This Book Is For Data scientists, big data experts, BI analysts, and data architects.

Over 90 hands-on recipes that will employ Nagios Core as the anchor of monitoring on your network About This Book Master the advanced configuration techniques of Nagios Core to model your network better by improving hosts, services, and contacts Filter and improve the notifications that Nagios Core sends in response to failed checks, which can greatly assist you when diagnosing problems Pull Nagios Core's data into a database to write clever custom reports of your own devise Who This Book Is For If you are a network or system administrator and are looking for instructions and examples on working with Nagios Core, then this book is for you. Some basic shell command-line experience is required, and some knowledge of scripting would be helpful when we discuss how plugins work. What You Will Learn Manage the configuration of Nagios Core with advanced techniques to achieve fine detail in your checks Find, install, and even write your own check plugins Filter notifications to send them to the right people or programs at the right time Work around difficult network accessibility issues and delegate checks to other machines Tweak a Nagios Core server to achieve both high performance and redundancy in case of disaster Process the results of checks performed by other machines to monitor backups and similar processes Extend Nagios Core to allow advanced scripting, reporting, and network visualization behavior In Detail Nagios Core is an open source monitoring framework suitable for any network that ensures both internal and customer-facing services are running correctly and manages notification and reporting behavior to diagnose and fix outages promptly. It allows very fine configuration of exactly when, where, what, and how to check network services to meet both the uptime goals of your network and systems team and the needs of your users. This book shows system and network administrators how to use Nagios Core to its fullest as a monitoring framework for checks on any kind of network services, from the smallest home network to much larger production multi-site services.

You will discover that Nagios Core is capable of doing much more than pinging a host or to see whether websites respond. The recipes in this book will demonstrate how to leverage Nagios Core's advanced configuration, scripting hooks, reports, data retrieval, and extensibility to integrate it with your existing systems, and to make it the rock-solid center of your network monitoring world. Style and approach Each chapter contains a set of step-by-step recipes to perform an example of a commonly performed task related to network administration. The book begins by focusing closely on the properties and configuration of Nagios Core itself, and gradually moves on to other pieces of software that can support, manage, and extend the system.

Real-world configurations and supporting materials enable you to deploy Nagios and integrate other tools on a step-by-step basis Simplifies deployment and installation by providing examples of real-world monitoring situations and explains how to configure, architect, and deploy EM solutions to address these situations Shows how to create your own Nagios plug-ins, to monitor devices for which Nagios doesn't provide plug-ins

FreeBSD and OpenBSD are increasingly gaining traction in educational institutions, non-profits, and corporations worldwide because they provide significant security advantages over Linux. Although a lot can be said for the robustness, clean organization, and stability of the BSD operating systems, security is one of the main reasons system administrators use these two platforms. There are plenty of books to help you get a FreeBSD or OpenBSD system off the ground, and all of them touch on security to some extent, usually dedicating a chapter to the subject. But, as security is commonly named as the key concern for today's system administrators, a single chapter on the subject can't provide the depth of information you need to keep your systems secure. FreeBSD and OpenBSD are rife with security "building blocks" that you can put to use, and Mastering FreeBSD and OpenBSD Security shows you how. Both operating systems have kernel options and filesystem features that go well beyond traditional Unix permissions and controls. This power and flexibility is valuable, but the colossal range of possibilities need to be tackled one step at a time. This book walks you through the installation of a hardened operating system, the installation and configuration of critical services, and ongoing maintenance of your FreeBSD and OpenBSD systems. Using an application-specific approach that builds on your existing knowledge, the book provides sound technical information on FreeBSD and Open-BSD security with plenty of real-world examples to help you configure and deploy a secure system. By imparting a solid technical foundation as well as practical know-how, it enables administrators to push their server's security to the next level. Even administrators in other environments--like Linux and Solaris--can find useful paradigms to emulate. Written by security professionals with two decades of operating system experience, Mastering FreeBSD and OpenBSD Security features broad and deep explanations of how how to secure your most critical systems. Where other books on BSD systems help you achieve

functionality, this book will help you more thoroughly secure your deployments.

Monitor your Elasticsearch cluster's health, and diagnose and solve its performance and reliability issues About This Book Understand common performance and reliability pitfalls in Elasticsearch Use popular monitoring tools such as Elasticsearch-head, BigDesk, Marvel, Kibana, and more This is a step-by-step guide with lots of case studies on solving real-world Elasticsearch cluster issues Who This Book Is For This book is for developers and system administrators who use Elasticsearch in a wide range of capacities. Prior knowledge of Elasticsearch and related technologies would be helpful, but is not necessary. What You Will Learn Explore your cluster with Elasticsearch-head and BigDesk Access the underlying data of the Elasticsearch monitoring plugins using the Elasticsearch API Analyze your cluster's performance with Marvel Troubleshoot some of the common performance and reliability issues that come up when using Elasticsearch Analyze a cluster's historical performance, and get to the bottom of and recover from system failures Use and install various other tools and plugins such as Kibana and Kopf, which is helpful to monitor Elasticsearch In Detail Elasticsearch is a distributed search server similar to Apache Solr with a focus on large datasets, a schema-less setup, and high availability. This schema-free architecture allows Elasticsearch to index and search unstructured content, making it perfectly suited for both small projects and large big data warehouses with petabytes of unstructured data. This book is your toolkit to teach you how to keep your cluster in good health, and show you how to diagnose and treat unexpected issues along the way. You will start by getting introduced to Elasticsearch, and look at some common performance issues that pop up when using the system. You will then see how to install and configure Elasticsearch and the Elasticsearch monitoring plugins. Then, you will proceed to install and use the Marvel dashboard to monitor Elasticsearch. You will find out how to troubleshoot some of the common performance and reliability issues that come up when using Elasticsearch. Finally, you will analyze your cluster's historical performance, and get to know how to get to the bottom of and recover from system failures. This book will guide you through several monitoring tools, and utilizes real-world cases and dilemmas faced when using Elasticsearch, showing you how to solve them simply, quickly, and cleanly. Style and approach This is a step-by-step guide to monitoring your Elasticsearch cluster and correcting performance issues. It is filled with lots of in-depth, real-world use-cases on solving different Elasticsearch cluster issues. Summary RabbitMQ in Action is a fast-paced run through building and managing scalable applications using the RabbitMQ messaging server. It starts by explaining how message queuing works, its history, and how RabbitMQ fits in. Then it shows you real-world examples you can apply to your own scalability and interoperability challenges. About the Technology There's a virtual switchboard at the core of most large applications where messages race between servers, programs, and services. RabbitMQ is an efficient and easy-to-deploy queue that handles this message traffic effortlessly

in all situations, from web startups to massive enterprise systems. About the Book RabbitMQ in Action teaches you to build and manage scalable applications in multiple languages using the RabbitMQ messaging server. It's a snap to get started. You'll learn how message queuing works and how RabbitMQ fits in. Then, you'll explore practical scalability and interoperability issues through many examples. By the end, you'll know how to make Rabbit run like a well-oiled machine in a 24 x 7 x 365 environment. Written for developers familiar with Python, PHP, Java, .NET, or any other modern programming language. No RabbitMQ experience required. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. What's Inside Learn fundamental messaging design patterns Use patterns for on-demand scalability Glue a PHP frontend to a backend written in anything Implement a PubSub-alerting service in 30 minutes flat Configure RabbitMQ's built-in clustering Monitor, manage, extend, and tune RabbitMQ =====?===== Table of Contents Pulling RabbitMQ out of the hat Understanding messaging Running and administering Rabbit Solving problems with Rabbit: coding and patterns Clustering and dealing with failure Writing code that survives failure Warrens and Shovels: failover and replication Administering RabbitMQ from the Web Controlling Rabbit with the REST API Monitoring: Houston, we have a problem Supercharging and securing your Rabbit Smart Rabbits: extending RabbitMQ

Get to grips with a new technology, understand what it is and what it can do for you, and then get to work with the most important features and tasks. A concise guide, written in an easy-to-follow format. Instant Nagios Starter is an outstanding resource for system engineers, administrators and developers, with a basic understanding of the Linux command line. Readers should have access to a test system or virtual machine with Linux installed to follow the given examples. Learn and monitor your entire IT infrastructure to ensure your systems, applications, services, and business function effectively. About This Book Packed with tips, tricks and illustrations, the book will explain the configuration and monitoring concepts in a simplified manner Experience the scalability and flexibility of Nagios in a very practical and easy-to-understand approach. Unleash the power of Nagios Core and Nagios XI 5 to monitor and secure your infrastructure with ease. Who This Book Is For This book is targeted at System Administrators, both, who have no prior knowledge of Nagios as well as readers experienced with it. It not only covers the basics of Nagios but also the advanced features. What You Will Learn Set up and use the built-in Nagios web interface Upskill the additional interfaces available for Nagios to monitor your IT infrastructure Learn how to perform various checks using both, Nagios standard plugins and third-party plugins Explore the working of notifications and events in Nagios Familiarize yourself with SNMP and use it for monitoring devices such as routers, switches, modems and printers Discover how can be Nagios can be customized and tailored to your needs Get to know more about the enterprise version of Nagios, Nagios XI In Detail Nagios, a powerful and widely used IT monitoring and management software for problem -solving. It detects problems related to your organizations infrastructure and helps in resolving the issue before it impacts the business. Following the success of the previous

edition, this book will continue to help you monitor the status of network devices and also notify the system administrators of network problems. Starting with the fundamentals, the book will teach you how to install and configure Nagios for your environment. The book helps you learn how to end downtimes, adding comments and generating reports using the built-in Web interface of Nagios. Moving on, you will be introduced to the third-party web interfaces and applications for checking the status and report specific information. As you progress further in Learning Nagios, you will focus on the standard set of Nagios plugins and also focus on teach you how to efficiently manage large configurations and using templates. Once you are up to speed with this, you will get to know the concept and working of notifications and events in Nagios. The book will then uncover the concept of passive check and shows how to use NRDP (Nagios Remote Data Processor). The focus then shifts to how Nagios checks can be run on remote machines and SNMP (Simple Network Management Protocol) can be used from Nagios. Lastly, the book will demonstrate how to extend Nagios by creating custom check commands, custom ways of notifying users and showing how passive checks and NRDP can be used to integrate your solutions with Nagios. By the end of the book, you will be a competent system administrator who could monitor mid-size businesses or even large scale enterprises. Style and approach This will be a practical learning guide for system administrators which will teach them everything about Nagios along with implementing it for your organization and then ending with securing it.

We can all be Linux experts, provided we invest the time in learning the craft of Linux administration. Pro Linux System Administration makes it easy for small- to medium-sized businesses to enter the world of zero-cost software running on Linux and covers all the distros you might want to use, including Red Hat, Ubuntu, Debian, and CentOS. Authors, and systems infrastructure experts James Turnbull, Peter Lieverdink, and Dennis Matotek take a layered, component-based approach to open source business systems, while training system administrators as the builders of business infrastructure. If you want to implement a SOHO or SMB Linux infrastructure, Pro Linux System Administration clearly demonstrates everything you need. You'll find this book also provides a solid framework to move forward and expand your business and associated IT capabilities, and you'll benefit from the expertise and experienced guidance of the authors. Pro Linux System Administration covers An introduction to using Linux and free and open source software to cheaply and efficiently manage your business A layered model that allows your infrastructure to grow with your business Easy and simple-to-understand instructions including configurations, examples, and extensive real-world hints and tips

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.

Practical Linux Infrastructure teaches you how to use the best open source tools to build a new Linux infrastructure, or alter an existing infrastructure, to ensure it stands up to enterprise-level needs. Each chapter covers a key area of implementation, with clear examples and step-by-step instructions. Using this book, you'll understand why scale matters, and what considerations you need to make. You'll see how to switch to using Google Cloud Platform for your hosted solution, how to use KVM for your virtualization, how to use Git, Postfix, and MySQL for your version control, email, and database, and how to use Puppet for your configuration management. For enterprise-level fault tolerance you'll use Apache, and for load balancing and high availability, you'll use HAProxy and Keepalived. For trend analysis you'll learn how to use Cacti, and for notification you'll use Nagios. You'll also learn how to utilize BIND to implement DNS, how to use DHCP (Dynamic Host Configuration Protocol), and how to setup remote access for your infrastructure using VPN and Iptables. You will finish by looking at the various tools you will need to troubleshoot issues that may occur with your hosted infrastructure. This includes how to use CPU, network, disk and memory management tools such as top, netstat, iostat and vmstat. Author Syed Ali is a senior site reliability engineering manager, who has extensive experience with virtualization and Linux cloud based infrastructure. His previous experience as an entrepreneur in infrastructure computing offers him deep insight into how a business can leverage the power of Linux to their advantage. He brings his expert knowledge to this book to teach others how to perfect their Linux environments. Become a Linux infrastructure pro with Practical Linux Infrastructure today.

This book will introduce Nagios to readers who are interested in monitoring their systems. All the concepts in the book are explained in a simplified manner, presented in an easy-to-understand language with lots of tips, tricks, and illustrations. This book

is great for system administrators interested in using Nagios to monitor their systems. It will also help professionals who have already worked with earlier versions of Nagios to understand the new features of Nagios 4 and provides usable solutions to real-life problems related to Nagios administration. To effectively use this book, system administration knowledge is required. If you want to create your own plug-ins, knowledge of scripting languages like Perl, shell and Python is expected.

Nagios 3 Enterprise Network Monitoring Including Plug-Ins and Hardware Devices Elsevier

Implement successful private clouds with OpenStack Key Features Gain hands-on experience in designing a private cloud for all infrastructures Create a robust virtual environment for your organization Design, implement and deploy an OpenStack-based cloud based on the Queens release Book Description Over the past six years, hundreds of organizations have successfully implemented Infrastructure as a Service (IaaS) platforms based on OpenStack. The huge amount of investment from these organizations, including industry giants such as IBM and HP, as well as open source leaders, such as Red Hat, Canonical, and SUSE, has led analysts to label OpenStack as the most important open source technology since the Linux operating system. Due to its ambitious scope, OpenStack is a complex and fast-evolving open source project that requires a diverse skill set to design and implement it. OpenStack for Architects leads you through the major decision points that you'll face while architecting an OpenStack private cloud for your organization. This book will address the recent changes made in the latest OpenStack release i.e Queens, and will also deal with advanced concepts such as containerization, NVF, and security. At each point, the authors offer you advice based on the experience they've gained from designing and leading successful OpenStack projects in a wide range of industries. Each chapter also includes lab material that gives you a chance to install and configure the technologies used to build production-quality OpenStack clouds. Most importantly, the book focuses on ensuring that your OpenStack project meets the needs of your organization, which will guarantee a successful rollout. What you will learn Learn the overall structure of an OpenStack deployment Craft an OpenStack deployment process which fits within your organization Apply Agile Development methodologies to engineer and operate OpenStack clouds Build a product roadmap for Infrastructure as a Service based on OpenStack Make use of containers to increase the manageability and resiliency of applications running in and on OpenStack. Use enterprise security guidelines for your OpenStack deployment Who this book is for OpenStack for Architects is for Cloud architects who are responsible to design and implement a private cloud with OpenStack. System engineers and enterprise architects will also find this book useful. Basic understanding of core OpenStack services, as well as some working experience of concepts, is recommended. High Performance MySQL is the definitive guide to building fast, reliable systems with MySQL. Written by noted experts with years of real-world experience building very large systems, this book covers every aspect of MySQL performance in detail, and focuses on robustness, security, and data integrity. High Performance MySQL teaches you advanced techniques in depth so you can bring out MySQL's full power. Learn how to design schemas, indexes, queries and advanced MySQL features for maximum performance, and get detailed guidance for tuning your MySQL server, operating system, and hardware to their fullest potential. You'll also learn practical, safe, high-performance ways to scale your applications with replication, load balancing, high availability,

and failover. This second edition is completely revised and greatly expanded, with deeper coverage in all areas. Major additions include: Emphasis throughout on both performance and reliability Thorough coverage of storage engines, including in-depth tuning and optimizations for the InnoDB storage engine Effects of new features in MySQL 5.0 and 5.1, including stored procedures, partitioned databases, triggers, and views A detailed discussion on how to build very large, highly scalable systems with MySQL New options for backups and replication Optimization of advanced querying features, such as full-text searches Four new appendices The book also includes chapters on benchmarking, profiling, backups, security, and tools and techniques to help you measure, monitor, and manage your MySQL installations.

Learn and monitor your entire IT infrastructure to ensure your systems, applications, services, and business function effectively About This Book *Packed with tips, tricks and illustrations, the book will explain the configuration and monitoring concepts in a simplified manner *Experience the scalability and flexibility of Nagios in a very practical and easy-to-understand approach. *Unleash the power of Nagios Core and Nagios XI 5 to monitor and secure your infrastructure with ease. Who This Book Is For This book is targeted at System Administrators, both, who have no prior knowledge of Nagios as well as readers experienced with it. It not only covers the basics of Nagios but also the advanced features. What You Will Learn *Set up and use the built-in Nagios web interface *Upskill the additional interfaces available for Nagios to monitor your IT infrastructure *Learn how to perform various checks using both, Nagios standard plugins and third-party plugins *Explore the working of notifications and events in Nagios *Familiarize yourself with SNMP and use it for monitoring devices such as routers, switches, modems and printers *Discover how can be Nagios can be customized and tailored to your needs *Get to know more about the enterprise version of Nagios, Nagios XI In Detail Nagios, a powerful and widely used IT monitoring and management software for problem-solving. It detects problems related to your organizations infrastructure and helps in resolving the issue before it impacts the business. Following the success of the previous edition, this book will continue to help you monitor the status of network devices and also notify the system administrators of network problems. Starting with the fundamentals, the book will teach you how to install and configure Nagios for your environment. The book helps you learn how to end downtimes, adding comments and generating reports using the built-in Web interface of Nagios. Moving on, you will be introduced to the third-party web interfaces and applications for checking the status and report specific information. As you progress further in Learning Nagios, you will focus on the standard set of Nagios plugins and also focus on teach you how to efficiently manage large configurations and using templates. Once you are up to speed with this, you will get to know the concept and working of notifications and events in Nagios. The book will then uncover the concept of passive check and shows how to use NRDP (Nagios Remote Data Processor). The focus then shifts to how Nagios checks can be run on remote machines and SNMP (Simple Network Management Protocol) can be used from Nagios. Lastly, the book will demonstrate how to extend Nagios by creating custom check commands, custom ways of notifying users and showing how passive checks and NRDP can be used to integrate your solutions with Nagios. By the end of the book, you will be a competent system administrator who could monitor mid-size businesses or even large scale enterprises.

This is the eBook of the printed book and may not include any media, website access codes, or print supplements that may come packaged with the bound book. The Comprehensive, Up-to-Date Apache Hadoop Administration Handbook and Reference “Sam Alapati has worked with production Hadoop clusters for six years. His unique depth of experience has enabled him to write the go-to resource for all administrators looking to spec, size, expand, and secure production Hadoop clusters of any size.” —Paul Dix, Series Editor In Expert Hadoop® Administration, leading Hadoop administrator Sam R. Alapati brings together authoritative knowledge for creating, configuring, securing, managing, and optimizing production Hadoop clusters in any environment. Drawing on his experience with large-scale Hadoop administration, Alapati integrates action-oriented advice with carefully researched explanations of both problems and solutions. He covers an unmatched range of topics and offers an unparalleled collection of realistic examples. Alapati demystifies complex Hadoop environments, helping you understand exactly what happens behind the scenes when you administer your cluster. You’ll gain unprecedented insight as you walk through building clusters from scratch and configuring high availability, performance, security, encryption, and other key attributes. The high-value administration skills you learn here will be indispensable no matter what Hadoop distribution you use or what Hadoop applications you run. Understand Hadoop’s architecture from an administrator’s standpoint Create simple and fully distributed clusters Run MapReduce and Spark applications in a Hadoop cluster Manage and protect Hadoop data and high availability Work with HDFS commands, file permissions, and storage management Move data, and use YARN to allocate resources and schedule jobs Manage job workflows with Oozie and Hue Secure, monitor, log, and optimize Hadoop Benchmark and troubleshoot Hadoop

Written by Ganglia designers and maintainers, this book shows you how to collect and visualize metrics from clusters, grids, and cloud infrastructures at any scale. Want to track CPU utilization from 50,000 hosts every ten seconds? Ganglia is just the tool you need, once you know how its main components work together. This hands-on book helps experienced system administrators take advantage of Ganglia 3.x. Learn how to extend the base set of metrics you collect, fetch current values, see aggregate views of metrics, and observe time-series trends in your data. You’ll also examine real-world case studies of Ganglia installs that feature challenging monitoring requirements. Determine whether Ganglia is a good fit for your environment Learn how Ganglia’s gmond and gmetad daemons build a metric collection overlay Plan for scalability early in your Ganglia deployment, with valuable tips and advice Take data visualization to a new level with gweb, Ganglia’s web frontend Write plugins to extend gmond’s metric-collection capability Troubleshoot issues you may encounter with a Ganglia installation Integrate Ganglia with the sFlow and Nagios monitoring systems Contributors include: Robert Alexander, Jeff Buchbinder, Frederiko Costa, Alex Dean, Dave Josephsen, Peter Phaal, and Daniel Pocock. Case study writers include: John Allspaw, Ramon Bastiaans, Adam Compton, Andrew Dibble, and Jonah Horowitz.

This book describes scientific results obtained by project partners and outcomes of research and development activities carried out within the Polish Infrastructure for Information Science Support in the European Research Space PL-Grid (PL-Grid 2011). Summary Mesos in Action introduces readers to the Apache Mesos cluster manager and the concept of application-centric

infrastructure. Filled with helpful figures and hands-on instructions, this book guides you from your first steps creating a highly-available Mesos cluster through deploying applications in production and writing native Mesos frameworks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Modern datacenters are complex environments, and when you throw Docker and other container-based systems into the mix, there's a great need to simplify. Mesos is an open source cluster management platform that transforms the whole datacenter into a single pool of compute, memory, and storage resources that you can allocate, automate, and scale as if you're working with a single supercomputer. About the Book Mesos in Action introduces readers to the Apache Mesos cluster manager and the concept of application-centric infrastructure. Filled with helpful figures and hands-on instructions, this book guides you from your first steps creating a highly-available Mesos cluster through deploying applications in production and writing native Mesos frameworks. You'll learn how to scale to thousands of nodes, while providing resource isolation between processes using Linux and Docker containers. You'll also learn practical techniques for deploying applications using popular key frameworks. What's Inside Spinning up your first Mesos cluster Scheduling, resource administration, and logging Deploying containerized applications with Marathon, Chronos, and Aurora Writing Mesos frameworks using Python About the Reader Readers need to be familiar with the core ideas of datacenter administration and need a basic knowledge of Python or a similar programming language. About the Author Roger Ignazio is an experienced systems engineer with a focus on distributed, fault-tolerant, and scalable infrastructure. He is currently a technical lead at Mesosphere. Table of Contents PART 1 HELLO, MESOS Introducing Mesos Managing datacenter resources with Mesos PART 2 CORE MESOS Setting up Mesos Mesos fundamentals Logging and debugging Mesos in production PART 3 RUNNING ON MESOS Deploying applications with MarathoN Managing scheduled tasks with Chronos Deploying applications and managing scheduled tasks with Aurora Developing a framework

Thoroughly updated and expanded! Includes new coverage on HANA, the cloud, and using SAP's applications! In just 24 sessions of one hour or less, you'll get up and running with the latest SAP technologies, applications, and solutions. Using a straightforward, step-by-step approach, each lesson strengthens your understanding of SAP from both a business and technical perspective, helping you gain practical mastery from the ground up on topics such as security, governance, validations, release management, SLA, and legal issues. Step-by-step instructions carefully walk you through the most common questions, issues, and tasks. Quizzes and exercises help you build and test your knowledge. Notes present interesting pieces of information. Tips offer advice or teach an easier way to do something. Cautions advise you about potential problems and help you steer clear of disaster. Learn how to... Understand SAP terminology, concepts, and solutions Install SAP on premises or in the cloud Master SAP's revamped user interface Discover how and when to use in-memory HANA databases Integrate SAP Software as a Service (SaaS) solutions such as Ariba, Successfactors, Fieldglass, and hybris Find resources at SAP's Service Marketplace, Developer Network, and Help Portal Avoid pitfalls in SAP project implementation, migration, and upgrades Discover how SAP fits with mobile devices, social media, big data, and the Internet of Things Start or accelerate your career working with SAP technologies

Download File PDF Monitoring With Nagios And Check Mk

[Copyright: f83d405c15ef2e79d28f74f7097035bf](#)