
Nagios Building Enterprise Grade Monitoring Infrastructures For Systems And Networks 2nd Edition

Practical Linux Infrastructure
Hands-On Enterprise Automation on Linux
Microservices from Theory to Practice: Creating Applications in IBM Bluemix Using the Microservices Approach
Nagios 3 Enterprise Network Monitoring
High Performance MySQL
Architecting the Digital Transformation
Monitoring with Ganglia
Nagios Core Administration Cookbook
Spring: Developing Java Applications for the Enterprise
Infrastructure Monitoring with Amazon CloudWatch
Innovative Technologies for Dependable OTS-Based Critical Systems
Web Operations
Building Microservices
Nagios 3 Enterprise Network Monitoring
Open Source Network Administration
Database Reliability Engineering
IBM Open Platform for DBaaS on IBM Power Systems
Security Monitoring
Learning Nagios 4
Data Lake for Enterprises
High Performance Scientific Computing Using Distributed Infrastructures
Distributed Applications and Interoperable Systems
Cloud Security and Privacy
Cloud Computing Advancements in Design, Implementation, and Technologies
Monitoring with Graphite
Building a Monitoring Infrastructure with Nagios
Building a National Distributed E-Infrastructure -- PL-Grid
Prometheus: Up & Running
Managing Digital
Datadog Cloud Monitoring Quick Start Guide
Practical Monitoring
Production-Ready Microservices
Hello, Startup
Nagios
Pro Puppet

Encyclopedia of E-Commerce Development, Implementation, and Management
Nagios, 2nd Edition
Linux Bible
Linux Advanced for SysAdmin
The Art of Monitoring

*Nagios Building Enterprise Grade
Monitoring Infrastructures For
Systems And Networks 2nd Edition*

Downloaded from archive.imba.com by
guest

GIOVANNA CLARA

Practical Linux Infrastructure "O'Reilly Media, Inc."

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.

Hands-On Enterprise Automation on Linux Packt Publishing Ltd
A practical guide to implementing your enterprise data lake using Lambda Architecture as the base
About This Book Build a full-fledged data lake for your organization with popular big data technologies using the Lambda architecture as the base
Delve into the big data technologies required to meet modern day business strategies A highly practical guide to implementing enterprise data lakes with lots of examples and real-world use-cases
Who This Book Is For Java developers and architects who would like to implement a data lake for their enterprise will find this book useful. If you want to get hands-on experience with the Lambda Architecture and big data technologies by implementing a practical solution using these technologies, this book will also help you.
What You Will Learn Build an enterprise-level data lake using the relevant big data technologies
Understand the core of the Lambda architecture and how to apply it in an enterprise
Learn the technical details around Sqoop and its functionalities
Integrate Kafka with Hadoop components to acquire enterprise data
Use flume with streaming technologies for stream-based processing
Understand stream-based processing with reference to Apache Spark Streaming
Incorporate Hadoop components and know the advantages they provide for enterprise data lakes
Build fast, streaming, and high-performance applications using Elasticsearch
Make your data ingestion process consistent across various data formats with configurability
Process your data to derive intelligence using machine learning algorithms
In Detail
The term "Data Lake" has recently emerged as a prominent term in the big data industry. Data scientists can make use of it in deriving meaningful insights that can be used by businesses to redefine or transform the way they operate. Lambda architecture is also emerging as one of the very eminent patterns in the big data landscape, as it not only helps to derive useful information from historical data but also correlates real-time data to enable business to take critical decisions. This book tries to bring these two important aspects — data lake and lambda

architecture—together. This book is divided into three main sections. The first introduces you to the concept of data lakes, the importance of data lakes in enterprises, and getting you up-to-speed with the Lambda architecture. The second section delves into the principal components of building a data lake using the Lambda architecture. It introduces you to popular big data technologies such as Apache Hadoop, Spark, Sqoop, Flume, and Elasticsearch. The third section is a highly practical demonstration of putting it all together, and shows you how an enterprise data lake can be implemented, along with several real-world use-cases. It also shows you how other peripheral components can be added to the lake to make it more efficient. By the end of this book, you will be able to choose the right big data technologies using the lambda architectural patterns to build your enterprise data lake.
Style and approach The book takes a pragmatic approach, showing ways to leverage big data technologies and lambda architecture to build an enterprise-level data lake.

Microservices from Theory to Practice: Creating Applications in IBM Bluemix Using the Microservices Approach "O'Reilly Media, Inc."

The infrastructure-as-code revolution in IT is also affecting database administration. With this practical book, developers, system administrators, and junior to mid-level DBAs will learn how the modern practice of site reliability engineering applies to the craft of database architecture and operations. Authors Laine Campbell and Charity Majors provide a framework for professionals looking to join the ranks of today's database reliability engineers (DBRE). You'll begin by exploring core operational concepts that DBREs need to master. Then you'll examine a wide range of database persistence options, including how to implement key technologies to provide resilient, scalable, and performant data storage and retrieval. With a firm foundation in database reliability engineering, you'll be ready to dive into the architecture and operations of any modern database. This book

covers: Service-level requirements and risk management Building and evolving an architecture for operational visibility Infrastructure engineering and infrastructure management How to facilitate the release management process Data storage, indexing, and replication Identifying datastore characteristics and best use cases Datastore architectural components and data-driven architectures

Nagios 3 Enterprise Network Monitoring IGI Global

A web application involves many specialists, but it takes people in web ops to ensure that everything works together throughout an application's lifetime. It's the expertise you need when your start-up gets an unexpected spike in web traffic, or when a new feature causes your mature application to fail. In this collection of essays and interviews, web veterans such as Theo Schlossnagle, Baron Schwartz, and Alistair Croll offer insights into this evolving field. You'll learn stories from the trenches--from builders of some of the biggest sites on the Web--on what's necessary to help a site thrive. Learn the skills needed in web operations, and why they're gained through experience rather than schooling Understand why it's important to gather metrics from both your application and infrastructure Consider common approaches to database architectures and the pitfalls that come with increasing scale Learn how to handle the human side of outages and degradations Find out how one company avoided disaster after a huge traffic deluge Discover what went wrong after a problem occurs, and how to prevent it from happening again Contributors include: John Allspaw Heather Champ Michael Christian Richard Cook Alistair Croll Patrick Debois Eric Florenzano Paul Hammond Justin Huff Adam Jacob Jacob Loomis Matt Massie Brian Moon Anoop Nagwani Sean Power Eric Ries Theo Schlossnagle Baron Schwartz Andrew Shafer

High Performance MySQL IBM Redbooks

Building upon fundamental Linux skills, this book is designed to elevate your expertise to handle intricate tasks crucial for senior system administrators. To begin with, we will revisit key administrative tasks, providing a solid foundation for the advanced topics to come. This includes managing files, directories, packages, system services, and user permissions to ensure you're up to speed. We then dive into the critical aspects of network management, wherein we learn to configure network interfaces, manage IP addresses and routing, set up and secure

network services like DHCP and DNS, and monitor network performance using tools such as Wireshark. We then move to explore iptables and firewalld, configure firewalls, implement AppArmor, and perform security audits with Lynis. The chapter also covers monitoring system logs for security breaches and using intrusion detection systems like Snort. Then, the chapter on Database Management demonstrates to work with databases in Linux, including installing and configuring PostgreSQL, designing databases, performing migrations, and securing database systems. This chapter equips you with the skills to manage database operations efficiently. The next chapter teaches you to monitor CPU and memory usage, track network performance, and analyze system logs using tools like Nagios and Zabbix. You will also learn to set up alerts and notifications to maintain optimal system performance. Later, we focus on automating routine tasks using shell scripting, cron, anacron, awk, and sed. You will learn to write scripts for system audits, resource management, and backup and recovery. Finally, the book covers large-scale deployments, Kubernetes, cluster management, load balancing, and kernel customization. Key Learnings Get the most out of AppArmor, firewalld, and iptables to boost up security. Take advantage of Nagios, Zabbix, and Wireshark to keep your systems and networks running smoothly. Get PostgreSQL up and running, migrate databases, and automate routine tasks; all while keeping databases secure. Resolve complex issues and maintain system health and uptime with troubleshooting skills. Optimize resource utilization, set up scalability, and ensure high availability by managing clusters. Improve your service's performance and dependability with sophisticated load balancing strategies. Personalize Linux's kernel in terms of speed, security, and hardware compatibility. Automate complex tasks with shell scripting, cron, and anacron. Table of Content Up and Running with System Administration Essentials Managing Networks Security and Monitoring Database Management System Health Monitoring Automation and Scripting Advanced System Administration

[Architecting the Digital Transformation](#) Springer Science & Business Media

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

Monitoring with Ganglia Packt Publishing Ltd

A comprehensive guide to rolling out Datadog to monitor infrastructure and applications running in both cloud and datacenter environments Key Features Learn Datadog to proactively monitor your infrastructure and cloud services Use Datadog as a platform for aggregating monitoring efforts in your organization Leverage Datadog's alerting service to implement on-call and site reliability engineering (SRE) processes Book Description Datadog is an essential cloud monitoring and operational analytics tool which enables the monitoring of servers, virtual machines, containers, databases, third-party tools, and application services. IT and DevOps teams can easily leverage Datadog to monitor infrastructure and cloud services, and this book will show you how. The book starts by describing basic monitoring concepts and types of monitoring that are rolled out in a large-scale IT production engineering environment. Moving on, the book covers how standard monitoring features are implemented on the Datadog platform and how they can be rolled out in a real-world production environment. As you advance, you'll discover how Datadog is integrated with popular software components that are used to build cloud platforms. The book also provides details on how to use monitoring standards such as Java Management Extensions (JMX) and StatsD to extend the Datadog platform. Finally, you'll get to grips with monitoring fundamentals, learn how monitoring can be rolled out using Datadog proactively, and find out how to extend and customize the Datadog platform. By the end of this Datadog book, you will have gained the skills needed to monitor your cloud infrastructure and the software applications running on it using Datadog. What you will learn Understand monitoring fundamentals, including metrics, monitors, alerts, and thresholds Implement core monitoring requirements using Datadog features Explore Datadog's integration with cloud platforms and tools Extend Datadog using custom scripting and standards such as JMX and StatsD Discover how proactive monitoring can be rolled out using various Datadog features Understand how Datadog can be used to monitor microservices in both Docker and Kubernetes environments Get to grips with advanced Datadog features such as APM and Security Monitoring Who this book is for This book is for DevOps engineers, site reliability engineers (SREs), IT Production engineers, software developers and architects, cloud engineers, system administrators, and anyone looking to monitor and visualize their

infrastructure and applications with Datadog. Basic working knowledge of cloud and infrastructure is useful. Working experience of Linux distribution and some scripting knowledge is required to fully take advantage of the material provided in the book.

Nagios Core Administration Cookbook GitforGits

This book is the "Hello, World" tutorial for building products, technologies, and teams in a startup environment. It's based on the experiences of the author, Yevgeniy (Jim) Brikman, as well as interviews with programmers from some of the most successful startups of the last decade, including Google, Facebook, LinkedIn, Twitter, GitHub, Stripe, Instagram, AdMob, Pinterest, and many others. Hello, Startup is a practical, how-to guide that consists of three parts: Products, Technologies, and Teams. Although at its core, this is a book for programmers, by programmers, only Part II (Technologies) is significantly technical, while the rest should be accessible to technical and non-technical audiences alike. If you're at all interested in startups—whether you're a programmer at the beginning of your career, a seasoned developer bored with large company politics, or a manager looking to motivate your engineers—this book is for you.

Spring: Developing Java Applications for the Enterprise

"O'Reilly Media, Inc."

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.

Infrastructure Monitoring with Amazon CloudWatch Prentice Hall
This book constitutes the refereed proceedings of the 9th IFIP WG 6.1 International Conference on Distributed Applications and Interoperable Systems, DAIS 2009, held in Lisbon, Portugal, in June 2009. The DAIS conference was held as part of the federated event on Distributed Computing Techniques (DisCoTec), together

with the 11th International Conference on Coordination Models and Languages (Coordination 2009) and the IFIP WG 6.1 International Conference on Formal Techniques for Distributed Systems (FMOODS/FORTE 2009). The 12 revised full papers presented were carefully reviewed and selected from 32 submissions. The papers address service orientation, quality of service and service contract, business processes, Web services, service components, algorithms and protocols supporting dependability, fault tolerance, data replication, group communication, adaptive and collaborative systems, context awareness, model-driven development, middleware for ubiquitous computing and sensor networks, ad hoc network protocols, peer-to-peer systems, and overlays. They are organized in topical sections peer-to-peer networks, adhoc networks, dependability, and infrastructure and services.

Innovative Technologies for Dependable OTS-Based Critical Systems "O'Reilly Media, Inc."

This research-oriented book presents key contributions on architecting the digital transformation. It includes the following main sections covering 20 chapters: · Digital Transformation · Digital Business · Digital Architecture · Decision Support · Digital Applications Focusing on digital architectures for smart digital products and services, it is a valuable resource for researchers, doctoral students, postgraduates, graduates, undergraduates, academics and practitioners interested in digital transformation. Web Operations "O'Reilly Media, Inc."

One of the biggest challenges for organizations that have adopted microservice architecture is the lack of architectural, operational, and organizational standardization. After splitting a monolithic application or building a microservice ecosystem from scratch, many engineers are left wondering what's next. In this practical book, author Susan Fowler presents a set of microservice standards in depth, drawing from her experience standardizing over a thousand microservices at Uber. You'll learn how to design microservices that are stable, reliable, scalable, fault tolerant, performant, monitored, documented, and prepared for any catastrophe. Explore production-readiness standards, including: Stability and Reliability: develop, deploy, introduce, and deprecate microservices; protect against dependency failures Scalability and Performance: learn essential components for achieving greater microservice efficiency Fault Tolerance and

Catastrophe Preparedness: ensure availability by actively pushing microservices to fail in real time Monitoring: learn how to monitor, log, and display key metrics; establish alerting and on-call procedures Documentation and Understanding: mitigate tradeoffs that come with microservice adoption, including organizational sprawl and technical debt

Building Microservices James Turnbull

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.

Nagios 3 Enterprise Network Monitoring World Scientific

Leverage the power of Spring MVC, Spring Boot, Spring Cloud, and additional popular web frameworks. About This Book Discover key Spring Framework-related technology standards such as Spring core, Spring-AOP, Spring data access frameworks, and Spring testing to develop robust Java applications easily This course is packed with tips and tricks that demonstrate Industry best practices on developing a Spring-MVC-based application Learn how to efficiently build and implement microservices in Spring, and how to use Docker and Mesos to push the boundaries and explore new possibilities Who This Book Is For This course is

intended for Java developers interested in building enterprise-level applications with Spring Framework. Prior knowledge of Java programming and web development concepts (and a basic knowledge of XML) is expected. What You Will Learn Understand the architecture of Spring Framework and how to set up the key components of the Spring Application Development Environment Configure Spring Container and manage Spring beans using XML and Annotation Practice Spring AOP concepts such as Aspect, Advice, Pointcut, and Introduction Integrate bean validation and custom validation Use error handling and exception resolving Get to grips with REST-based web service development and Ajax Use Spring Boot to develop microservices Find out how to avoid common pitfalls when developing microservices Get familiar with end-to-end microservices written in Spring Framework and Spring Boot In Detail This carefully designed course aims to get you started with Spring, the most widely adopted Java framework, and then goes on to more advanced topics such as building microservices using Spring Boot within Spring. With additional coverage of popular web frameworks such as Struts, WebWork, Java Server Faces, Tapestry, Docker, and Mesos, you'll have all the skills and expertise you need to build great applications. Starting with the Spring Framework architecture and setting up the key components of the Spring Application Development Environment, you will learn how to configure Spring Container and manage Spring beans using XML and Annotation. Next, you will delve into Spring MVC, which will help you build flexible and loosely coupled web applications. You'll also get to grips with testing applications for reliability. Moving on, this course will help you implement the microservice architecture in Spring Framework, Spring Boot, and Spring Cloud. Written to the latest specifications of Spring, this book will help you build modern, Internet-scale Java applications in no time. This Learning Path combines some of the best that Packt has to offer in one complete, curated package. It includes content from the following Packt products: Learning Spring Application Development by Ravi Kant Soni Spring MVC Beginner's Guide - Second Edition by Amuthan Ganeshan Spring Microservices by Rajesh RV Style and approach This is a step-by-step guide for building a complete application and developing scalable microservices using Spring Framework, Spring Boot, and a set of Spring Cloud components [Open Source Network Administration](#) Packt Publishing Ltd

Pro Puppet is an in-depth guide to installing, using, and developing the popular configuration management tool Puppet. The book is a comprehensive follow-up to the previous title Pulling Strings with Puppet. Puppet provides a way to automate everything from user management to server configuration. You'll learn how to create Puppet recipes, extend Puppet, and use Facter to gather configuration data from your servers. 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!

Database Reliability Engineering Packt Publishing Ltd
A hands-on and introductory guide to the art of modern application and infrastructure monitoring and metrics. We start small and then build on what you learn to scale out to multi-site, multi-tier applications. The book is written for both developers and sysadmins. We focus on building monitored and measurable applications. We also use tools that are designed to handle the challenges of managing Cloud, containerised and distributed applications and infrastructure. In the book we'll deliver: * An introduction to monitoring, metrics and measurement. * A scalable framework for monitoring hosts (including Docker and containers), services and applications built on top of the Riemann event stream processor. * Graphing and metric storage using Graphite and Grafana. * Logging with Logstash. * A framework for high quality and useful notifications * Techniques for developing and building monitorable applications * A capstone that puts all the pieces together to monitor a multi-tier application.
IBM Open Platform for DBaaS on IBM Power Systems "O'Reilly Media, Inc."

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

Security Monitoring IGI Global

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).

Learning Nagios 4 Springer Nature

The demand for large-scale dependable, systems, such as Air Traffic Management, industrial plants and space systems, is attracting efforts of many world-leading European companies and SMEs in the area, and is expected to increase in the near future. The adoption of Off-The-Shelf (OTS) items plays a key role in such a scenario. OTS items allow mastering complexity and reducing costs and time-to-market; however, achieving these goals by ensuring dependability requirements at the same time is challenging. CRITICAL STEP project establishes a strategic collaboration between academic and industrial partners, and proposes a framework to support the development of dependable, OTS-based, critical systems. The book introduces methods and tools adopted by the critical systems industry, and surveys key achievements of the CRITICAL STEP project along four directions: fault injection tools, V&V of critical systems, runtime monitoring and evaluation techniques, and security assessment.

Data Lake for Enterprises "O'Reilly Media, Inc."

Microservices is an architectural style in which large, complex software applications are composed of one or more smaller services. Each of these microservices focuses on completing one task that represents a small business capability. These microservices can be developed in any programming language. They communicate with each other using language-neutral protocols, such as Representational State Transfer (REST), or messaging applications, such as IBM® MQ Light. This IBM Redbooks® publication gives a broad understanding of this increasingly popular architectural style, and provides some real-life examples of how you can develop applications using the microservices approach with IBM Bluemix™. The source code for all of these sample scenarios can be found on GitHub

(<https://github.com/>). The book also presents some case studies from IBM products. We explain the architectural decisions made,

our experiences, and lessons learned when redesigning these products using the microservices approach. Information technology (IT) professionals interested in learning about

microservices and how to develop or redesign an application in Bluemix using microservices can benefit from this book.

Related with Nagios Building Enterprise Grade Monitoring Infrastructures For Systems And Networks 2nd Edition:

- Printable Teenage Anger Management Worksheets : [click here](#)