

---

# Apache Nifi Overview Linux Foundation

---

Handle Data-Driven Challenges in an Enterprise Big Data Lake

Learning Spark

Practical Enterprise Data Lake Insights

Enterprise Data Warehouse Optimization with Hadoop on IBM Power Systems Servers

The Success of Open Source

Monitoring with Ganglia

The Cathedral & the Bazaar

Work with massive datasets to design data models and automate data pipelines using Python

AI and Big Data on IBM Power Systems Servers

Deep Learning

A Hands-On Guide to Building Robust and Scalable Event-Driven Applications with Code Examples in Java

Advanced Platform Development with Kubernetes

Architecting Modern Data Platforms

Real-Time Data and Stream Processing at Scale

Enabling Data Management, the Internet of Things, Blockchain, and Machine Learning

Lightning-Fast Big Data Analysis

An Introduction

Apache Solr Enterprise Search Server - Third Edition

Enterprise Integration Patterns

Accounting and Business Ethics

Big Data Analytics

Quality of Information and Communications Technology

Agile Data Science

Elegance, Evolution, and a Few Fearless Hacks

Effective Kafka

Molly Fyde and the Parsona Rescue

The Architecture of Open Source Applications

New Designs Using Apache Kafka and Mapr Streams

Introducing .NET 4.5

Kafka: The Definitive Guide

A Guide to Enterprise Hadoop at Scale

Data Pipelines with Apache Airflow

Building Big Data Applications

Charting Your Strategy for Next-Generation Business Analytics

Enterprise Integration Patterns

Learning Spark

Advances in Information and Communication

---

## **CUMMINGS NATHANIEL**

---

Handle Data-Driven Challenges in an Enterprise Big Data Lake Addison-Wesley Professional

Leverage Kubernetes for the rapid adoption of emerging technologies. Kubernetes is the future of enterprise platform development and has become the most popular, and often considered the most robust, container orchestration system available today. This book focuses on platforming technologies that power the Internet of Things, Blockchain, Machine Learning, and the many layers of data and application management supporting them. Advanced Platform Development with Kubernetes takes you through the process of building platforms with these in-demand capabilities. You'll progress through the development of Serverless, CICD integration, data processing pipelines, event queues, distributed query engines, modern data warehouses, data lakes, distributed object storage, indexing and analytics, data routing and transformation, query engines, and data science/machine learning environments. You'll also see how to implement and tie together numerous essential and trending technologies including: Kafka, NiFi, Airflow, Hive, Keycloak, Cassandra, MySQL, Zookeeper, Mosquitto, Elasticsearch, Logstash, Kibana, Presto, Mino, OpenFaaS, and Ethereum. The book uses Golang and Python to demonstrate the development integration of custom container and Serverless functions, including

interaction with the Kubernetes API. The exercises throughout teach Kubernetes through the lens of platform development, expressing the power and flexibility of Kubernetes with clear and pragmatic examples. Discover why Kubernetes is an excellent choice for any individual or organization looking to embark on developing a successful data and application platform. What You'll Learn Configure and install Kubernetes and k3s on vendor-neutral platforms, including generic virtual machines and bare metal Implement an integrated development toolchain for continuous integration and deployment Use data pipelines with MQTT, NiFi, Logstash, Kafka and Elasticsearch Install a serverless platform with OpenFaaS Explore blockchain network capabilities with Ethereum Support a multi-tenant data science platform and web IDE with JupyterHub, MLflow and Seldon Core Build a hybrid cluster, securely bridging on-premise and cloud-based Kubernetes nodes Who This Book Is For System and software architects, full-stack developers, programmers, and DevOps engineers with some experience building and using containers. This book also targets readers who have started with Kubernetes and need to progress from a basic understanding of the technology and "Hello World" example to more productive, career-building projects. Learning Spark "O'Reilly Media, Inc." Data in all domains is getting bigger. How can you work with it efficiently? Recently updated for Spark 1.3, this book introduces Apache Spark, the open source cluster computing system that makes data analytics fast to write and fast to run. With Spark, you can tackle

big datasets quickly through simple APIs in Python, Java, and Scala. This edition includes new information on Spark SQL, Spark Streaming, setup, and Maven coordinates. Written by the developers of Spark, this book will have data scientists and engineers up and running in no time. You'll learn how to express parallel jobs with just a few lines of code, and cover applications from simple batch jobs to stream processing and machine learning. Quickly dive into Spark capabilities such as distributed datasets, in-memory caching, and the interactive shell Leverage Spark's powerful built-in libraries, including Spark SQL, Spark Streaming, and MLlib Use one programming paradigm instead of mixing and matching tools like Hive, Hadoop, Mahout, and Storm Learn how to deploy interactive, batch, and streaming applications Connect to data sources including HDFS, Hive, JSON, and S3 Master advanced topics like data partitioning and shared variables [Practical Enterprise Data Lake Insights](#) "O'Reilly Media, Inc."

Build on your existing programming skills and upskill to professional-level C# programming. Summary In Code Like A Pro in C# you will learn: Unit testing and test-driven development Refactor a legacy .NET codebase Principles of clean code Essential backend architecture skills Query and manipulate databases with LINQ and Entity Framework Core Critical business applications worldwide are written in the versatile C# language and the powerful .NET platform, running on desktops, cloud systems, and Windows or Linux servers. Code Like a Pro in C# makes it easy to turn your existing abilities in C# or another OO language (such as Java) into practical C# mastery. There's no "Hello World" or Computer Science 101 basics—you'll

learn by refactoring an out-of-date legacy codebase, using new techniques, tools, and best practices to bring it up to modern C# standards. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology You know the basics, now get ready for the next step! Pro-quality C# code is efficient, clean, and fast. Whether you're building user-facing business applications or writing data-intensive backend services, the experience-based, practical techniques in this book will take your C# skills to a new level. About the book Code Like a Pro in C# teaches you to how write clean C# code that's suitable for enterprise applications. In this book, you'll refactor a legacy codebase by applying modern C# techniques. You'll explore tools like Entity Framework Core, design techniques like dependency injection, and key practices like testing and clean coding. It's a perfect path to upgrade your existing C# skills or shift from another OO language into C# and the .NET ecosystem. What's inside Unit testing and test-driven development Refactor a legacy .NET codebase Principles of clean code Query and manipulate databases with LINQ and Entity Framework Core About the reader For developers experienced with object-oriented programming. No C# experience required. About the author Jort Rodenburg is a software engineer who has taught numerous courses on getting up to speed with C# and .NET. Table of Contents PART 1 USING C# AND .NET 1 Introducing C# and .NET 2 .NET and how it compiles PART 2 THE EXISTING CODEBASE 3 How bad is this code? 4 Manage your unmanaged resources! PART 3 THE DATABASE ACCESS LAYER 5 Setting up a project

and database with Entity Framework Core PART 4 THE REPOSITORY LAYER 6 Test-driven development and dependency injection 7 Comparing objects 8 Stubbing, generics, and coupling 9 Extension methods, streams, and abstract classes PART 5 THE SERVICE LAYER 10 Reflection and mocks 11 Runtime type checking revisited and error handling 12 Using IAsyncEnumerable and yield return PART 6 THE CONTROLLER LAYER 13 Middleware, HTTP routing, and HTTP responses 14 JSON serialization/deserialization and custom model binding

### **Enterprise Data Warehouse**

#### **Optimization with Hadoop on IBM**

#### **Power Systems Servers IBM Redbooks**

This book is a comprehensive introduction to building data pipelines, that will have you moving and transforming data in no time. You'll learn how to build data pipelines, transform and clean data, and deliver it to provide value to users. You will learn to deploy production data pipelines that include logging, monitoring, and version control. [The Success of Open Source](#) O'Reilly Media

Flow Architectures"O'Reilly Media, Inc."

*Monitoring with Ganglia* Harvard

University Press

Data Pipelines with Apache Airflow teaches you the ins-and-outs of the Directed Acyclic Graphs (DAGs) that power Airflow, and how to write your own DAGs to meet the needs of your projects. With complete coverage of both foundational and lesser-known features, when you're done you'll be set to start using Airflow for seamless data pipeline development and management.

Pipelines can be challenging to manage, especially when your data has to flow through a collection of application

components, servers, and cloud services. Airflow lets you schedule, restart, and backfill pipelines, and its easy-to-use UI and workflows with Python scripting has users praising its incredible flexibility. Data Pipelines with Apache Airflow takes you through best practices for creating pipelines for multiple tasks, including data lakes, cloud deployments, and data science. Data Pipelines with Apache Airflow teaches you the ins-and-outs of the Directed Acyclic Graphs (DAGs) that power Airflow, and how to write your own DAGs to meet the needs of your projects. With complete coverage of both foundational and lesser-known features, when you're done you'll be set to start using Airflow for seamless data pipeline development and management.

Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

*The Cathedral & the Bazaar* Apress Building Big Data Applications helps data managers and their organizations make the most of unstructured data with an existing data warehouse. It provides readers with what they need to know to make sense of how Big Data fits into the world of Data Warehousing. Readers will learn about infrastructure options and integration and come away with a solid understanding on how to leverage various architectures for integration. The book includes a wide range of use cases that will help data managers visualize reference architectures in the context of specific industries (healthcare, big oil, transportation, software, etc.). Explores various ways to leverage Big Data by effectively integrating it into the data warehouse Includes real-world case studies which clearly demonstrate Big Data technologies Provides insights on how to optimize current data warehouse

infrastructure and integrate newer infrastructure matching data processing workloads and requirements

**Work with massive datasets to design data models and automate data pipelines using Python** Apress

The Complete Guide to Data Science with Hadoop—For Technical Professionals, Businesspeople, and Students Demand is soaring for professionals who can solve real data science problems with Hadoop and Spark. Practical Data Science with Hadoop® and Spark is your complete guide to doing just that. Drawing on immense experience with Hadoop and big data, three leading experts bring together everything you need: high-level concepts, deep-dive techniques, real-world use cases, practical applications, and hands-on tutorials. The authors introduce the essentials of data science and the modern Hadoop ecosystem, explaining how Hadoop and Spark have evolved into an effective platform for solving data science problems at scale. In addition to comprehensive application coverage, the authors also provide useful guidance on the important steps of data ingestion, data munging, and visualization. Once the groundwork is in place, the authors focus on specific applications, including machine learning, predictive modeling for sentiment analysis, clustering for document analysis, anomaly detection, and natural language processing (NLP). This guide provides a strong technical foundation for those who want to do practical data science, and also presents business-driven guidance on how to apply Hadoop and Spark to optimize ROI of data science initiatives. Learn What data science is, how it has evolved, and how to plan a data science career How data volume, variety, and velocity shape data

science use cases Hadoop and its ecosystem, including HDFS, MapReduce, YARN, and Spark Data importation with Hive and Spark Data quality, preprocessing, preparation, and modeling Visualization: surfacing insights from huge data sets Machine learning: classification, regression, clustering, and anomaly detection Algorithms and Hadoop tools for predictive modeling Cluster analysis and similarity functions Large-scale anomaly detection NLP: applying data science to human language

*AI and Big Data on IBM Power Systems Servers* Apress

Would you like to use a consistent visual notation for drawing integration solutions? "Look inside the front cover." Do you want to harness the power of asynchronous systems without getting caught in the pitfalls? "See "Thinking Asynchronously" in the Introduction." Do you want to know which style of application integration is best for your purposes? "See Chapter 2, Integration Styles." Do you want to learn techniques for processing messages concurrently? "See Chapter 10, Competing Consumers and Message Dispatcher." Do you want to learn how you can track asynchronous messages as they flow across distributed systems? "See Chapter 11, Message History and Message Store." Do you want to understand how a system designed using integration patterns can be implemented using Java Web services, .NET message queuing, and a TIBCO-based publish-subscribe architecture? "See Chapter 9, Interlude: Composed Messaging." Utilizing years of practical experience, seasoned experts Gregor Hohpe and Bobby Woolf show how asynchronous messaging has proven to be the best strategy for enterprise integration success. However,

building and deploying messaging solutions presents a number of problems for developers. "Enterprise Integration Patterns" provides an invaluable catalog of sixty-five patterns, with real-world solutions that demonstrate the formidable of messaging and help you to design effective messaging solutions for your enterprise. The authors also include examples covering a variety of different integration technologies, such as JMS, MSMQ, TIBCO ActiveEnterprise, Microsoft BizTalk, SOAP, and XSL. A case study describing a bond trading system illustrates the patterns in practice, and the book offers a look at emerging standards, as well as insights into what the future of enterprise integration might hold. This book provides a consistent vocabulary and visual notation framework to describe large-scale integration solutions across many technologies. It also explores in detail the advantages and limitations of asynchronous messaging architectures. The authors present practical advice on designing code that connects an application to a messaging system, and provide extensive information to help you determine when to send a message, how to route it to the proper destination, and how to monitor the health of a messaging system. If you want to know how to manage, monitor, and maintain a messaging system once it is in use, get this book. 0321200683B09122003

#### Flow Architectures

This book is for developers who want to learn how to get the most out of Solr in their applications, whether you are new to the field, have used Solr but don't know everything, or simply want a good reference. It would be helpful to have some familiarity with basic programming concepts, but no prior experience is required.

#### *Deep Learning* "O'Reilly Media, Inc."

The software architecture landscape has evolved dramatically over the past decade. Microservices have displaced monoliths. Data and applications are increasingly becoming distributed and decentralised. But composing disparate systems is a hard problem. More recently, software practitioners have been rapidly converging on event-driven architecture as a sustainable way of dealing with complexity - integrating systems without increasing their coupling. In *Effective Kafka*, Emil Koutanov explores the fundamentals of Event-Driven Architecture - using Apache Kafka - the world's most popular and supported open-source event streaming platform. You'll learn:

- The fundamentals of event-driven architecture and event streaming platforms-
- The background and rationale behind Apache Kafka, its numerous potential uses and applications-
- The architecture and core concepts - the underlying software components, partitioning and parallelism, load-balancing, record ordering and consistency modes-
- Installation of Kafka and related tooling - using standalone deployments, clusters, and containerised deployments with Docker-
- Using CLI tools to interact with and administer Kafka classes, as well as publishing data and browsing topics-
- Using third-party web-based tools for monitoring a cluster and gaining insights into the event streams-
- Building stream processing applications in Java 11 using off-the-shelf client libraries-
- Patterns and best-practice for organising the application architecture, with emphasis on maintainability and testability of the resulting code-
- The numerous gotchas that lurk in Kafka's client and broker configuration, and how to counter them-
- Theoretical background on distributed



and concurrent computing, exploring factors affecting their liveness and safety- Best-practices for running multi-tenanted clusters across diverse engineering teams, how teams collaborate to build complex systems at scale and equitably share the cluster with the aid of quotas- Operational aspects of running Kafka clusters at scale, performance tuning and methods for optimising network and storage utilisation- All aspects of Kafka security - including network segregation, encryption, certificates, authentication and authorization. The coverage is progressively delivered and carefully aimed at giving you a journey-like experience into becoming proficient with Apache Kafka and Event-Driven Architecture. The goal is to get you designing and building applications. And by the conclusion of this book, you will be a confident practitioner and a Kafka evangelist within your organisation - wielding the knowledge necessary to teach others.

*A Hands-On Guide to Building Robust and Scalable Event-Driven Applications with Code Examples in Java* "O'Reilly Media, Inc."

Growing up an orphan in the Milky Way hasn't been easy, especially as a teenage girl in the Naval Academy. Unfortunately for Molly Fyde, things are about to get worse. Just as she's finding her place amongst the boys, her unfair expulsion from the Academy takes away the only two things that truly matter: flying in space and her training partner, Cole. Sent off to a normal school, she feels destined for a dull, unspectacular future. Then, a marvelous discovery changes everything: Her father's old starship, missing for a decade, turns up halfway across the galaxy. Its retrieval launches Molly and Cole on the

adventure of a lifetime, one that will have lasting consequences for themselves and billions of others. What starts off as a simple quest to reconnect with her past, ends up forging a new future. And the forgotten family she hoped to uncover is replaced by a new one she never foresaw: a band of alien misfits and runaways... The crew of the starship Parsona.

[Advanced Platform Development with Kubernetes](#) Packt Publishing Ltd

Data warehouses were developed for many good reasons, such as providing quick query and reporting for business operations, and business performance. However, over the years, due to the explosion of applications and data volume, many existing data warehouses have become difficult to manage. Extract, Transform, and Load (ETL) processes are taking longer, missing their allocated batch windows. In addition, data types that are required for business analysis have expanded from structured data to unstructured data. The Apache open source Hadoop platform provides a great alternative for solving these problems. IBM® has committed to open source since the early years of open Linux. IBM and Hortonworks together are committed to Apache open source software more than any other company. IBM Power Systems™ servers are built with open technologies and are designed for mission-critical data applications. Power Systems servers use technology from the OpenPOWER Foundation, an open technology infrastructure that uses the IBM POWER® architecture to help meet the evolving needs of big data applications. The combination of Power Systems with Hortonworks Data Platform (HDP) provides users with a highly efficient platform that provides

leadership performance for big data workloads such as Hadoop and Spark. This IBM Redpaper™ publication provides details about Enterprise Data Warehouse (EDW) optimization with Hadoop on Power Systems. Many people know Power Systems from the IBM AIX® platform, but might not be familiar with IBM PowerLinux™, so part of this paper provides a Power Systems overview. A quick introduction to Hadoop is provided for those not familiar with the topic. Details of HDP on Power Reference architecture are included that will help both software architects and infrastructure architects understand the design. In the optimization chapter, we describe various topics: traditional EDW offload, sizing guidelines, performance tuning, IBM Elastic Storage™ Server (ESS) for data-intensive workload, IBM Big SQL as the common structured query language (SQL) engine for Hadoop platform, and tools that are available on Power Systems that are related to EDW optimization. We also dedicate some pages to the analytics components (IBM Data Science Experience (IBM DSX) and IBM Spectrum™ Conductor for Spark workload) for the Hadoop infrastructure.

### **Architecting Modern Data Platforms** Springer Nature

Mining big data requires a deep investment in people and time. How can you be sure you're building the right models? With this hands-on book, you'll learn a flexible toolset and methodology for building effective analytics applications with Hadoop. Using lightweight tools such as Python, Apache Pig, and the D3.js library, your team will create an agile environment for exploring data, starting with an example application to mine your own email inboxes. You'll learn an iterative approach that enables you to quickly

change the kind of analysis you're doing, depending on what the data is telling you. All example code in this book is available as working Heroku apps. Create analytics applications by using the agile big data development methodology Build value from your data in a series of agile sprints, using the data-value stack Gain insight by using several data structures to extract multiple features from a single dataset Visualize data with charts, and expose different aspects through interactive reports Use historical data to predict the future, and translate predictions into action Get feedback from users after each sprint to keep your project on track [Real-Time Data and Stream Processing at Scale](#) "O'Reilly Media, Inc." Beschrijving van vijftientig open source applicaties.

### **Enabling Data Management, the Internet of Things, Blockchain, and Machine Learning** Penguin

Learn advanced analytical techniques and leverage existing tool kits to make your analytic applications more powerful, precise, and efficient. This book provides the right combination of architecture, design, and implementation information to create analytical systems that go beyond the basics of classification, clustering, and recommendation. Pro Hadoop Data Analytics emphasizes best practices to ensure coherent, efficient development. A complete example system will be developed using standard third-party components that consist of the tool kits, libraries, visualization and reporting code, as well as support glue to provide a working and extensible end-to-end system. The book also highlights the importance of end-to-end, flexible, configurable, high-performance data pipeline systems with analytical



components as well as appropriate visualization results. You'll discover the importance of mix-and-match or hybrid systems, using different analytical components in one application. This hybrid approach will be prominent in the examples. What You'll Learn Build big data analytic systems with the Hadoop ecosystem Use libraries, tool kits, and algorithms to make development easier and more effective Apply metrics to measure performance and efficiency of components and systems Connect to standard relational databases, noSQL data sources, and more Follow case studies with example components to create your own systems Who This Book Is For Software engineers, architects, and data scientists with an interest in the design and implementation of big data analytical systems using Hadoop, the Hadoop ecosystem, and other associated technologies.

*Lightning-Fast Big Data Analysis* Simon and Schuster

More and more data-driven companies are looking to adopt stream processing and streaming analytics. With this concise ebook, you'll learn best practices for designing a reliable architecture that supports this emerging big-data paradigm. Authors Ted Dunning and Ellen Friedman (Real World Hadoop) help you explore some of the best technologies to handle stream processing and analytics, with a focus on the upstream queuing or message-passing layer. To illustrate the effectiveness of these technologies, this book also includes specific use cases. Ideal for developers and non-technical people alike, this book describes: Key elements in good design for streaming analytics, focusing on the essential characteristics of the messaging layer New messaging technologies,

including Apache Kafka and MapR Streams, with links to sample code Technology choices for streaming analytics: Apache Spark Streaming, Apache Flink, Apache Storm, and Apache Apex How stream-based architectures are helpful to support microservices Specific use cases such as fraud detection and geo-distributed data streams Ted Dunning is Chief Applications Architect at MapR Technologies, and active in the open source community. He currently serves as VP for Incubator at the Apache Foundation, as a champion and mentor for a large number of projects, and as committer and PMC member of the Apache ZooKeeper and Drill projects. Ted is on Twitter as @ted\_dunning. Ellen Friedman, a committer for the Apache Drill and Apache Mahout projects, is a solutions consultant and well-known speaker and author, currently writing mainly about big data topics. With a PhD in Biochemistry, she has years of experience as a research scientist and has written about a variety of technical topics. Ellen is on Twitter as @Ellen\_Friedman.

**An Introduction** Routledge

This book presents high-quality research on the concepts and developments in the field of information and communication technologies, and their applications. It features 134 rigorously selected papers (including 10 poster papers) from the Future of Information and Communication Conference 2020 (FICC 2020), held in San Francisco, USA, from March 5 to 6, 2020, addressing state-of-the-art intelligent methods and techniques for solving real-world problems along with a vision of future research. Discussing various aspects of communication, data science, ambient intelligence, networking, computing,

security and Internet of Things, the book offers researchers, scientists, industrial engineers and students valuable insights into the current research and next generation information science and communication technologies.

*Apache Solr Enterprise Search Server - Third Edition* "O'Reilly Media, Inc."

Get up to speed with Apache Drill, an extensible distributed SQL query engine that reads massive datasets in many popular file formats such as Parquet, JSON, and CSV. Drill reads data in HDFS or in cloud-native storage such as S3 and works with Hive metastores along with distributed databases such as HBase, MongoDB, and relational databases. Drill works everywhere: on your laptop or in your largest cluster. In this practical book, Drill committers Charles Givre and Paul Rogers show analysts and data scientists how to query and analyze raw data using this powerful tool. Data scientists today spend about 80% of their time just gathering and cleaning data. With this book, you'll learn how Drill helps you analyze data more effectively to drive down time to insight. Use Drill to clean, prepare, and summarize delimited data for further analysis Query file types including logfiles, Parquet, JSON, and other complex formats Query Hadoop, relational databases, MongoDB, and Kafka with standard SQL Connect to Drill programmatically using a variety of languages Use Drill even with challenging or ambiguous file formats Perform sophisticated analysis by extending Drill's functionality with user-defined functions Facilitate data analysis for network security, image metadata, and machine learning

Enterprise Integration Patterns Apress Enhance your skills in building scalable infrastructure for your cloud-based

applications Key Features Learn to design a scalable architecture by building continuous integration (CI) pipelines with Kubernetes Get an in-depth understanding of role-based access control (RBAC), continuous deployment (CD), and observability Monitor a Kubernetes cluster with Prometheus and Grafana Book Description Kubernetes is among the most popular open-source platforms for automating the deployment, scaling, and operations of application containers across clusters of hosts, providing a container-centric infrastructure. Hands-On Microservices with Kubernetes starts by providing you with in-depth insights into the synergy between Kubernetes and microservices. You will learn how to use Delinkcious, which will serve as a live lab throughout the book to help you understand microservices and Kubernetes concepts in the context of a real-world application. Next, you will get up to speed with setting up a CI/CD pipeline and configuring microservices using Kubernetes ConfigMaps. As you cover later chapters, you will gain hands-on experience in securing microservices, and implementing REST, gRPC APIs, and a Delinkcious data store. In addition to this, you'll explore the Nuclio project, run a serverless task on Kubernetes, and manage and implement data-intensive tests. Toward the concluding chapters, you'll deploy microservices on Kubernetes and learn to maintain a well-monitored system. Finally, you'll discover the importance of service meshes and how to incorporate Istio into the Delinkcious cluster. By the end of this book, you'll have gained the skills you need to implement microservices on Kubernetes with the help of effective tools and best practices. What you will learn Understand the synergy between

Kubernetes and microservices Create a complete CI/CD pipeline for your microservices on Kubernetes Develop microservices on Kubernetes with the Go kit framework using best practices Manage and monitor your system using Kubernetes and open-source tools Expose your services through REST and gRPC APIs Implement and deploy serverless functions as a service Externalize authentication, authorization and traffic shaping using a service mesh Run a Kubernetes cluster in the cloud on

Google Kubernetes Engine Who this book is for This book is for developers, DevOps engineers, or anyone who wants to develop large-scale microservice-based systems on top of Kubernetes. If you are looking to use Kubernetes on live production projects or want to migrate existing systems to a modern containerized microservices system, then this book is for you. Coding skills, together with some knowledge of Docker, Kubernetes, and cloud concepts will be useful.

Related with Apache Nifi Overview Linux Foundation:

- Brad Treliving Trade History : [click here](#)