

# Download Ebook Release It Design And Deploy Production Ready Software Pragmatic Programmers Michael T Nygard Read Pdf Free

Release It! Release It! Engineering Production-Grade Shiny Apps Deploy Machine Learning Models to Production Deploy Machine Learning Models to Production Continuous Delivery in Java Production-Ready Microservices Programming AWS Lambda Site Reliability Engineering OpenShift for Developers Real-Time Phoenix Continuous Delivery in Java Deploying Rails Continuous Delivery Production Ready OpenStack - Recipes for Successful Environments Deployment with Docker Kubernetes in Production Best Practices Building Machine Learning Pipelines High Performance Drupal Deploying Node.js Production-Ready Applied Deep Learning Deployment Fundamentals, Vol. 6 Crystal Clear Cloud Foundry: The Definitive Guide React 17 Design Patterns and Best Practices Learn Amazon SageMaker Trino: The Definitive Guide Terraform: Up & Running Ansible: Up and Running Accelerate Deep Learning for Coders with fastai and PyTorch Practical Docker with Python Engineering MLOps Cloud Native Python Learning Serverless Microservices in Action Microservice Patterns and Best Practices Practical Machine Learning with AWS Deep Learning Systems Kubernetes for Developers

**Accelerate** Dec 19 2021 Winner of the Shingo Publication Award Accelerate your organization to win in the marketplace. How can we apply technology to drive business value? For years, we've been told that the performance of software delivery teams doesn't matter—that it can't provide a competitive advantage to our companies. Through four years of groundbreaking research to include data collected from the State of DevOps reports conducted with Puppet, Dr. Nicole Forsgren, Jez Humble, and Gene Kim set out to find a way to measure software delivery performance—and what drives it?—using rigorous statistical methods. This book presents both the findings and the science behind that research, making the information accessible for readers to apply in their own organizations. Readers will discover how to measure the performance of their teams, and what capabilities they should invest in to drive higher performance. This book is ideal for management at every level.

**Site Reliability Engineering** Oct 09 2023 The overwhelming majority of a software system's lifespan is spent in use, not in design or implementation. So, why does conventional wisdom insist that software engineers focus primarily on the design and development of large-scale computing systems? In this collection of essays and articles, key members of Google's Site Reliability Team explain how and why their commitment to the entire lifecycle has enabled the company to successfully build, deploy, monitor, and maintain some of the largest software systems in the world. You'll learn the principles and practices that enable Google engineers to make systems more scalable, reliable, and efficient—lessons directly applicable to your organization. This book is divided into four sections: Introduction—Learn what site reliability engineering is and why it differs from conventional IT industry practices Principles—Examine the patterns, behaviors, and areas of concern that influence the work of a site reliability engineer (SRE) Practices—Understand the theory and practice of an SRE's day-to-day work: building and operating large distributed computing systems Management—Explore Google's best practices for training, communication, and meetings that your organization can use

**Continuous Delivery** May 04 2023 Winner of the 2011 Jolt Excellence Award! Getting software released to users is often a painful, risky, and time-consuming process. This groundbreaking new book sets out the principles and technical practices that enable rapid, incremental delivery of high quality, valuable new functionality to users. Through automation of the build, deployment, and testing process, and improved collaboration between developers, testers, and operations, delivery teams can get changes released in a matter of hours—sometimes even minutes—no matter what the size of a project or the complexity of its code base. Jez Humble and David Farley begin by presenting the foundations of a rapid, reliable, low-risk delivery process. Next, they introduce the “deployment pipeline,” an automated process for managing all changes, from check-in to release. Finally, they discuss the “ecosystem” needed to support continuous delivery, from infrastructure, data and configuration management to governance. The authors introduce state-of-the-art techniques, including automated infrastructure management and data migration, and the use of virtualization. For each, they review key issues, identify best practices, and demonstrate how to mitigate risks. Coverage includes • Automating all facets of building, integrating, testing, and deploying software • Implementing deployment pipelines at team and organizational levels • Improving collaboration between developers, testers, and operations • Developing features incrementally on large and distributed teams • Implementing an effective configuration management strategy • Automating acceptance testing, from analysis to implementation • Testing capacity and other non-functional requirements • Implementing continuous deployment and zero-downtime releases • Managing infrastructure, data, components and dependencies • Navigating risk management, compliance, and auditing Whether you're a developer, systems administrator, tester, or manager, this book will help your organization move from idea to release faster than ever—so you can deliver value to your business rapidly and reliably.

**Deep Learning Systems** Mar 10 2021 This book describes deep learning systems: the algorithms, compilers, and processor components to efficiently train and deploy deep learning models for commercial applications. The exponential growth in computational power is slowing at a time when the amount of compute consumed by state-of-the-art deep learning (DL) workloads is rapidly growing. Model size, serving latency, and power constraints are a significant challenge in the deployment of DL models for many applications. Therefore, it is imperative to codesign algorithms, compilers, and hardware to accelerate advances in this field with holistic system-level and algorithm solutions that improve performance, power, and efficiency. Advancing DL systems generally involves three types of engineers: (1) data scientists that utilize and develop DL algorithms in partnership with domain experts, such as medical, economic, or climate scientists; (2) hardware designers that develop specialized hardware to accelerate the components in the DL models; and (3) performance and compiler engineers that optimize software to run more efficiently on a given hardware. Hardware engineers should be aware of the characteristics and components of production and academic models likely to be adopted by industry to guide design decisions impacting future hardware. Data scientists should be aware of deployment platform constraints when designing models. Performance engineers should support optimizations across diverse models, libraries, and hardware targets. The purpose of this book is to provide a solid understanding of (1) the design, training, and applications of DL algorithms in industry; (2) the compiler techniques to map deep learning code to hardware targets; and (3) the critical hardware features that accelerate DL systems. This book aims to facilitate co-innovation for the advancement of DL systems. It is written for engineers working in one or more of these areas who seek to understand the entire system stack in order to better collaborate with engineers working in other parts of the system stack. The book details advancements and adoption of DL models in industry, explains the training and deployment process, describes the essential hardware architectural features needed for today's and future models, and details advances in DL compilers to efficiently execute algorithms across various hardware targets. Unique in this book is the holistic exposition of the entire DL system stack, the emphasis on commercial applications, and the practical techniques to design models and accelerate their performance. The author is fortunate to work with hardware, software, data scientist, and research teams across many high-technology companies with hyperscale data centers. These companies employ many of the examples and methods provided throughout the book.

**Deploying Node.js** Oct 29 2022 About This Book Gain comprehensive information on scaling Node.js both vertically and horizontally in the cloud and across virtual machines Learn how to use Gulp, Mocha, Chai, PhantomJS, Git, Browserify, Docker, and other tools to construct a simple, fast, and intelligent build-and-deploy system Utilize the easy-to-understand examples to explore data caching strategies, application monitoring tips and tricks, and other professional techniques that are useful for maintaining lightweight, fast, and robust applications Who This Book Is For If you are an intermediate or advanced developer deploying your Node.js applications, then this book is for you. If you have already built a Node application or module and want to take your knowledge to the next level, this book will help you find your way. What You Will Learn Learn the strengths of Node.js and how to optimize your code to

make it ready for deployment into production Use Docker and Vagrant to create many virtual instances of your Node.js applications Deploy Node.js repositories to the cloud using Heroku, OpenShift, and DigitalOcean Utilize native Node.js modules, or Nginx, to load balance your applications Optimize your Node runtime by tweaking V8, managing memory intelligently, and using Redis to manage state data Discover how to use Gulp, Browserify, npm, Mocha, Chai, Sinon, and other tools to simplify your build/test/release process Manage your production deployments with Git, Jenkins, and Ansible In Detail Node.js is a popular choice for teams that need to design, build, test, deploy, maintain, and monitor large-scale distributed systems. Starting with a detailed overview of the Node.js architecture, this book covers topics that will help in application development, testing, deployment, and maintenance. You will learn about concurrency, event loops, callbacks and streams. Furthermore, step-by-step instructions on deploying applications to providers such as DigitalOcean and Heroku will be provided, including information on setting up load balancers and proxies. Message queues and other techniques for managing state and session data at scale will also be covered. A series of examples on deploying your Node.js applications in production environments are provided, including a discussion on setting up continuous deployment and integration for your team. Popular tools for testing, deploying, building, and monitoring Node.js applications are covered, helping you get up and running quickly.

**Continuous Delivery in Java** Jul 06 2023 Continuous delivery adds enormous value to the business and the entire software delivery lifecycle, but adopting this practice means mastering new skills typically outside of a developer's comfort zone. In this practical book, Daniel Bryant and Abraham Marín-Pérez provide guidance to help experienced Java developers master skills such as architectural design, automated quality assurance, and application packaging and deployment on a variety of platforms. Not only will you learn how to create a comprehensive build pipeline for continually delivering effective software, but you'll also explore how Java application architecture and deployment platforms have affected the way we rapidly and safely deliver new software to production environments. Get advice for beginning or completing your migration to continuous delivery Design architecture to enable the continuous delivery of Java applications Build application artifacts including fat JARs, virtual machine images, and operating system container (Docker) images Use continuous integration tooling like Jenkins, PMD, and find-sec-bugs to automate code quality checks Create a comprehensive build pipeline and design software to separate the deploy and release processes Explore why functional and system quality attribute testing is vital from development to delivery Learn how to effectively build and test applications locally and observe your system while it runs in production

**Release It!** May 16 2024 A single dramatic software failure can cost a company millions of dollars - but can be avoided with simple changes to design and architecture. This new edition of the best-selling industry standard shows you how to create systems that run longer, with fewer failures, and recover better when bad things happen. New coverage includes DevOps, microservices, and cloud-native architecture. Stability antipatterns have grown to include systemic problems in large-scale systems. This is a must-have pragmatic guide to engineering for production systems. If you're a software developer, and you don't want to get alerts every night for the rest of your life, help is here. With a combination of case studies about huge losses - lost revenue, lost reputation, lost time, lost opportunity - and practical, down-to-earth advice that was all gained through painful experience, this book helps you avoid the pitfalls that cost companies millions of dollars in downtime and reputation. Eighty percent of project life-cycle cost is in production, yet few books address this topic. This updated edition deals with the production of today's systems - larger, more complex, and heavily virtualized - and includes information on chaos engineering, the discipline of applying randomness and deliberate stress to reveal systematic problems. Build systems that survive the real world, avoid downtime, implement zero-downtime upgrades and continuous delivery, and make cloud-native applications resilient. Examine ways to architect, design, and build software - particularly distributed systems - that stands up to the typhoon winds of a flash mob, a Slashdotting, or a link on Reddit. Take a hard look at software that failed the test and find ways to make sure your software survives. To skip the pain and get the experience...get this book.

**Practical Docker with Python** Oct 17 2021 Learn the key differences between containers and virtual machines. Adopting a project based approach, this book introduces you to a simple Python application to be developed and containerized with Docker. After an introduction to Containers and Docker you'll be guided through Docker installation and configuration. You'll also learn basic functions and commands used in Docker by running a simple container using Docker commands. The book then moves on to developing a Python based Messaging Bot using required libraries and virtual environment where you'll add Docker Volumes to your project, ensuring your container data is safe. You'll create a database container and link your project to it and finally, bring up the Bot-associated database all at once with Docker Compose. What You'll Learn Build, run, and distribute Docker containers Develop a Python App and containerize it Use Dockerfile to run the Python App Define and run multi-container applications with Docker Compose Work with persisting data generated by and used by Docker containers Who This Book Is For Intermediate developers/DevOps practitioners who are looking to improve their build and release workflow by containerizing applications

**Crystal Clear** Jul 26 2022 Carefully researched over ten years and eagerly anticipated by the agile community, Crystal Clear: A Human-Powered Methodology for Small Teams is a lucid and practical introduction to running a successful agile project in your organization. Each chapter illuminates a different important aspect of orchestrating agile projects. Highlights include Attention to the essential human and communication aspects of successful projects Case studies, examples, principles, strategies, techniques, and guiding properties Samples of work products from real-world projects instead of blank templates and toy problems Top strategies used by software teams that excel in delivering quality code in a timely fashion Detailed introduction to emerging best-practice techniques, such as Blitz Planning, Project 360o, and the essential Reflection Workshop Question-and-answer with the author about how he arrived at these recommendations, including where they fit with CMMI, ISO, RUP, XP, and other methodologies A detailed case study, including an ISO auditor's analysis of the project Perhaps the most important contribution this book offers is the Seven Properties of Successful Projects. The author has studied successful agile projects and identified common traits they share. These properties lead your project to success; conversely, their absence endangers your project.

**High Performance Drupal** Nov 29 2022 How can you help your Drupal website continue to perform at the highest level as it grows to meet demand? This comprehensive guide provides best practices, examples, and in-depth explanations for solving several performance and scalability issues. You'll learn how to apply coding and infrastructure techniques to Drupal internals, application performance, databases, web servers, and performance analysis. Covering Drupal versions 7 and 8, this book is the ideal reference for everything from site deployment to implementing specific technologies such as Varnish, memcache, or Solr. If you have a basic understanding of Drupal and the Linux-Apache-MySQL-PHP (LAMP) stack, you're ready to get started. Establish a performance baseline and define goals for improvement Optimize your website's code and front-end performance Get best and worst practices for customizing Drupal core functionality Apply infrastructure design techniques to launch or expand a site Use tools to configure, monitor, and optimize MySQL performance Employ alternative storage and backend search options as your site grows Tune your web servers through httpd and PHP configuration Monitor services and perform load tests to catch problems before they become critical

**Deploy Machine Learning Models to Production** Mar 14 2024 Build and deploy machine learning and deep learning models in production with end-to-end examples. This book begins with a focus on the machine learning model deployment process and its related challenges. Next, it covers the process of building and deploying machine learning models using different web frameworks such as Flask and Streamlit. A chapter on Docker follows and covers how to package and containerize machine learning models. The book also illustrates how to build and train machine learning and deep learning models at scale using Kubernetes. The book is a good starting point for people who want to move to the next level of machine learning by taking pre-built models and deploying them into production. It also offers guidance to those who want to move beyond Jupyter notebooks to training models at scale on cloud environments. All the code presented in the book is available in the form of Python scripts for you to try the examples and extend them in interesting ways. What You Will Learn Build, train, and deploy machine learning models at scale using Kubernetes Containerize any kind of machine learning model and run it on any platform using Docker Deploy machine learning and deep learning models using Flask and Streamlit frameworks Who This Book Is For Data engineers, data scientists, analysts, and machine learning and deep learning engineers

**Engineering Production-Grade Shiny Apps** Apr 15 2024 From the Reviews "[This book] contains an excellent blend of both Shiny-specific topics ... and practical advice from software development that fits in nicely with Shiny apps. You will find many nuggets of wisdom sprinkled throughout these chapters...." Eric Nantz, Host of the R-Podcast and the Shiny Developer Series (from the Foreword) "[This] book is a gradual and pleasant invitation to the production-ready shiny apps world. It ...exposes a comprehensive and robust workflow powered by the {golem} package. [It] fills the not yet covered gap between shiny app development and deployment in such a thrilling way that it may be read in one sitting.... In the industry world, where processes robustness is a key toward productivity, this book will indubitably have a tremendous impact." David Granjon, Sr. Expert Data Science, Novartis Presented in full

color, Engineering Production-Grade Shiny Apps helps people build production-grade shiny applications, by providing advice, tools, and a methodology to work on web applications with R. This book starts with an overview of the challenges which arise from any big web application project: organizing work, thinking about the user interface, the challenges of teamwork and the production environment. Then, it moves to a step-by-step methodology that goes from the idea to the end application. Each part of this process will cover in detail a series of tools and methods to use while building production-ready shiny applications. Finally, the book will end with a series of approaches and advice about optimizations for production. Features Focused on practical matters: This book does not cover Shiny concepts, but practical tools and methodologies to use for production. Based on experience: This book is a formalization of several years of experience building Shiny applications. Original content: This book presents new methodologies and tooling, not just a review of what already exists. Engineering Production-Grade Shiny Apps covers medium to advanced content about Shiny, so it will help people that are already familiar with building apps with Shiny, and who want to go one step further.

**Ansible: Up and Running** Jan 20 2022 Among the many configuration management tools available, Ansible has some distinct advantages—it's minimal in nature, you don't need to install anything on your nodes, and it has an easy learning curve. This practical guide shows you how to be productive with this tool quickly, whether you're a developer deploying code to production or a system administrator looking for a better automation solution. Author Lorin Hochstein shows you how to write playbooks (Ansible's configuration management scripts), manage remote servers, and explore the tool's real power: built-in declarative modules. You'll discover that Ansible has the functionality you need and the simplicity you desire. Understand how Ansible differs from other configuration management systems Use the YAML file format to write your own playbooks Learn Ansible's support for variables and facts Work with a complete example to deploy a non-trivial application Use roles to simplify and reuse playbooks Make playbooks run faster with ssh multiplexing, pipelining, and parallelism Deploy applications to Amazon EC2 and other cloud platforms Use Ansible to create Docker images and deploy Docker containers

**Kubernetes for Developers** Feb 06 2021 A clear and practical beginner's guide that shows you just how easy it can be to make the switch to Kubernetes! Kubernetes for Developers reveals practical and painless methods for deploying your apps on Kubernetes—even for small-to-medium sized applications! You'll learn how to migrate your existing apps onto Kubernetes without a rebuild, and implement modern cloud native architectures that can handle your future growth. Inside, you'll learn how to: Containerize a web application with Docker Host a containerized app on Kubernetes with a public cloud service Save money and improve performance with cloud native technologies Make your deployments reliable and fault tolerant Prepare your deployments to scale without a redesign Monitor, debug and tune application deployments on Kubernetes Designed for busy working developers, this hands-on guide helps your first steps into Kubernetes using the powerful Google Kubernetes Engine (GKE) service. Learn how the GKE's powerful automation tools can perform automatic checks and scaling, giving you more time to spend developing great applications. You'll soon see that you don't need to incur huge costs or have the manpower of an enterprise organization to get a productivity boost from Kubernetes! About the technology Modern software needs to perform at scale while effectively handling load balancing, state and security. Kubernetes makes these tasks easier and more reliable for apps of any size. This book, written especially for software developers creating applications that run on Kubernetes, shows you exactly how to address these and other important issues. About the book Kubernetes for Developers covers everything you need to know to containerize and deploy an app on Kubernetes from the developer's perspective. You'll start by creating a small application you can run on a cloud-based Kubernetes cluster. Then, you'll systematically explore best practices for stable long-term deployment, including scaling, capacity planning, and resource optimization. What's inside Deploying reliable web applications using automated operations Scaling up without an application redesign Monitoring, debugging, and tuning workloads About the reader For developers familiar with building or deploying web applications. No Docker or Kubernetes experience required. About the author William Denniss is a product manager at Google working on Google Kubernetes Engine. Table of Contents PART 1 Getting started with Kubernetes 1 Kubernetes for application deployment 2 Containerizing apps 3 Deploying to Kubernetes 4 Automated operations 5 Resource management PART 2 Going to production 6 Scaling up 7 Internal services and load balancing 8 Node feature selection 9 Stateful applications 10 Background processing 11 GitOps: Configuration as code 12 Securing Kubernetes

**Terraform: Up & Running** Feb 18 2022 Terraform has become a key player in the DevOps world for defining, launching, and managing infrastructure as code (IaC) across a variety of cloud and virtualization platforms, including AWS, Google Cloud, Azure, and more. This hands-on second edition, expanded and thoroughly updated for Terraform version 0.12 and beyond, shows you the fastest way to get up and running. Gruntwork cofounder Yevgeniy (Jim) Brikman walks you through code examples that demonstrate Terraform's simple, declarative programming language for deploying and managing infrastructure with a few commands. Veteran sysadmins, DevOps engineers, and novice developers will quickly go from Terraform basics to running a full stack that can support a massive amount of traffic and a large team of developers. Explore changes from Terraform 0.9 through 0.12, including backends, workspaces, and first-class expressions Learn how to write production-grade Terraform modules Dive into manual and automated testing for Terraform code Compare Terraform to Chef, Puppet, Ansible, CloudFormation, and Salt Stack Deploy server clusters, load balancers, and databases Use Terraform to manage the state of your infrastructure Create reusable infrastructure with Terraform modules Use advanced Terraform syntax to achieve zero-downtime deployment Microservices in Action Jun 12 2021 "The one [and only] book on implementing microservices with a real-world, cover-to-cover example you can relate to." - Christian Bach, Swiss Re Microservices in Action is a practical book about building and deploying microservice-based applications. Written for developers and architects with a solid grasp of service-oriented development, it tackles the challenge of putting microservices into production. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Invest your time in designing great applications, improving infrastructure, and making the most out of your dev teams. Microservices are easier to write, scale, and maintain than traditional enterprise applications because they're built as a system of independent components. Master a few important new patterns and processes, and you'll be ready to develop, deploy, and run production-quality microservices. About the Book Microservices in Action teaches you how to write and maintain microservice-based applications. Created with day-to-day development in mind, this informative guide immerses you in real-world use cases from design to deployment. You'll discover how microservices enable an efficient continuous delivery pipeline, and explore examples using Kubernetes, Docker, and Google Container Engine. What's inside An overview of microservice architecture Building a delivery pipeline Best practices for designing multi-service transactions and queries Deploying with containers Monitoring your microservices About the Reader Written for intermediate developers familiar with enterprise architecture and cloud platforms like AWS and GCP. About the Author Morgan Bruce and Paulo A. Pereira are experienced engineering leaders. They work daily with microservices in a production environment, using the techniques detailed in this book. Table of Contents Designing and running microservices Microservices at SimpleBank Architecture of a microservice application Designing new features Transactions and queries in microservices Designing reliable services Building a reusable microservice framework Deploying microservices Deployment with containers and schedulers Building a delivery pipeline for microservices Building a monitoring system Using logs and traces to understand behavior Building microservice teams PART 1 - The lay of the land PART 2 - Design PART 3 - Deployment PART 4 - Observability and ownership

*Learn Amazon SageMaker* Apr 22 2022 Quickly build and deploy machine learning models without managing infrastructure, and improve productivity using Amazon SageMaker's capabilities such as Amazon SageMaker Studio, Autopilot, Experiments, Debugger, and Model Monitor Key Features Build, train, and deploy machine learning models quickly using Amazon SageMaker Analyze, detect, and receive alerts relating to various business problems using machine learning algorithms and techniques Improve productivity by training and fine-tuning machine learning models in production Book Description Amazon SageMaker enables you to quickly build, train, and deploy machine learning (ML) models at scale, without managing any infrastructure. It helps you focus on the ML problem at hand and deploy high-quality models by removing the heavy lifting typically involved in each step of the ML process. This book is a comprehensive guide for data scientists and ML developers who want to learn the ins and outs of Amazon SageMaker. You'll understand how to use various modules of SageMaker as a single toolset to solve the challenges faced in ML. As you progress, you'll cover features such as AutoML, built-in algorithms and frameworks, and the option for writing your own code and algorithms to build ML models. Later, the book will show you how to integrate Amazon SageMaker with popular deep learning libraries such as TensorFlow and PyTorch to increase the capabilities of existing models. You'll also learn to get the models to production faster with minimum effort and at a lower cost. Finally, you'll explore how to use Amazon SageMaker Debugger to analyze, detect, and highlight problems to understand the current model state and improve model accuracy. By the end of this Amazon book, you'll be able to use Amazon SageMaker on the full spectrum of ML workflows, from experimentation, training, and monitoring to scaling, deployment, and automation. What you will learn Create and automate end-to-end machine learning

workflows on Amazon Web Services (AWS) Become well-versed with data annotation and preparation techniques Use AutoML features to build and train machine learning models with AutoPilot Create models using built-in algorithms and frameworks and your own code Train computer vision and NLP models using real-world examples Cover training techniques for scaling, model optimization, model debugging, and cost optimization Automate deployment tasks in a variety of configurations using SDK and several automation tools Who this book is for This book is for software engineers, machine learning developers, data scientists, and AWS users who are new to using Amazon SageMaker and want to build high-quality machine learning models without worrying about infrastructure. Knowledge of AWS basics is required to grasp the concepts covered in this book more effectively. Some understanding of machine learning concepts and the Python programming language will also be beneficial.

**Cloud Native Python** Aug 15 2021 Build cloud native applications in Python About This Book This is the only reliable resource that showcases the tools and techniques you need build robust and resilient cloud native applications in Python Learn how to architect your application on both, the AWS and Azure clouds for high availability Assess, monitor, and troubleshoot your applications in the cloud Who This Book Is For This book is ideal for developers with a basic knowledge of Python who want to learn to build, test, and scale their Python-based applications. No prior experience of writing microservices in Python is required. What You Will Learn Get to know “the way of the cloud”, including why developing good cloud software is fundamentally about mindset and discipline Know what microservices are and how to design them Create reactive applications in the cloud with third-party messaging providers Build massive-scale, user-friendly GUIs with React and Flux Secure cloud-based web applications: the do's, don'ts, and options Plan cloud apps that support continuous delivery and deployment In Detail Businesses today are evolving so rapidly that having their own infrastructure to support their expansion is not feasible. As a result, they have been resorting to the elasticity of the cloud to provide a platform to build and deploy their highly scalable applications. This book will be the one stop for you to learn all about building cloud-native architectures in Python. It will begin by introducing you to cloud-native architecture and will help break it down for you. Then you'll learn how to build microservices in Python using REST APIs in an event driven approach and you will build the web layer. Next, you'll learn about Interacting data services and building Web views with React, after which we will take a detailed look at application security and performance. Then, you'll also learn how to Dockerize your services. And finally, you'll learn how to deploy the application on the AWS and Azure platforms. We will end the book by discussing some concepts and techniques around troubleshooting problems that might occur with your applications after you've deployed them. This book will teach you how to craft applications that are built as small standard units, using all the proven best practices and avoiding the usual traps. It's a practical book: we're going to build everything using Python 3 and its amazing tooling ecosystem. The book will take you on a journey, the destination of which, is the creation of a complete Python application based on microservices over the cloud platform Style and approach Filled with examples, this book takes a step-by-step approach to teach you each and every configuration you need to make your application highly available and fault tolerant.

**Engineering MLOps** Sep 15 2021 Get up and running with machine learning life cycle management and implement MLOps in your organization Key Features Become well-versed with MLOps techniques to monitor the quality of machine learning models in production Explore a monitoring framework for ML models in production and learn about end-to-end traceability for deployed models Perform CI/CD to automate new implementations in ML pipelines Book Description Engineering MLOps presents comprehensive insights into MLOps coupled with real-world examples in Azure to help you to write programs, train robust and scalable ML models, and build ML pipelines to train and deploy models securely in production. The book begins by familiarizing you with the MLOps workflow so you can start writing programs to train ML models. Then you'll then move on to explore options for serializing and packaging ML models post-training to deploy them to facilitate machine learning inference, model interoperability, and end-to-end model traceability. You'll learn how to build ML pipelines, continuous integration and continuous delivery (CI/CD) pipelines, and monitor pipelines to systematically build, deploy, monitor, and govern ML solutions for businesses and industries. Finally, you'll apply the knowledge you've gained to build real-world projects. By the end of this ML book, you'll have a 360-degree view of MLOps and be ready to implement MLOps in your organization. What you will learn Formulate data governance strategies and pipelines for ML training and deployment Get to grips with implementing ML pipelines, CI/CD pipelines, and ML monitoring pipelines Design a robust and scalable microservice and API for test and production environments Curate your custom CD processes for related use cases and organizations Monitor ML models, including monitoring data drift, model drift, and application performance Build and maintain automated ML systems Who this book is for This MLOps book is for data scientists, software engineers, DevOps engineers, machine learning engineers, and business and technology leaders who want to build, deploy, and maintain ML systems in production using MLOps principles and techniques. Basic knowledge of machine learning is necessary to get started with this book.

**Trino: The Definitive Guide** Mar 22 2022 Perform fast interactive analytics against different data sources using the Trino high-performance distributed SQL query engine. With this practical guide, you'll learn how to conduct analytics on data where it lives, whether it's Hive, Cassandra, a relational database, or a proprietary data store. Analysts, software engineers, and production engineers will learn how to manage, use, and even develop with Trino. Initially developed by Facebook, open source Trino is now used by Netflix, Airbnb, LinkedIn, Twitter, Uber, and many other companies. Matt Fuller, Manfred Moser, and Martin Traverso show you how a single Trino query can combine data from multiple sources to allow for analytics across your entire organization. Get started: Explore Trino's use cases and learn about tools that will help you connect to Trino and query data Go deeper: Learn Trino's internal workings, including how to connect to and query data sources with support for SQL statements, operators, functions, and more Put Trino in production: Secure Trino, monitor workloads, tune queries, and connect more applications; learn how other organizations apply Trino

**Deployment with Docker** Mar 02 2023 A practical guide to rapidly and efficiently mastering Docker containers, along with tips and tricks learned in the field. About This Book Use Docker containers, horizontal node scaling, modern orchestration tools (Docker Swarm, Kubernetes, and Mesos) and Continuous Integration/Continuous Delivery to manage your infrastructure. Increase service density by turning often-idle machines into hosts for numerous Docker services. Learn what it takes to build a true container infrastructure that is scalable, reliable, and resilient in the face of increased complexities from using container infrastructures. Find out how to identify, debug, and mitigate most real-world, undocumented issues when deploying your own Docker infrastructure. Learn tips and tricks of the trade from existing Docker infrastructures running in production environments. Who This Book Is For This book is aimed at system administrators, developers, DevOps engineers, and software engineers who want to get concrete, hands-on experience deploying multi-tier web applications and containerized microservices using Docker. This book is also for anyone who has worked on deploying services in some fashion and wants to take their small-scale setups to the next level (or simply to learn more about the process). What You Will Learn Set up a working development environment and create a simple web service to demonstrate the basics Learn how to make your service more usable by adding a database and an app server to process logic Add resilience to your services by learning how to horizontally scale with a few containers on a single node Master layering isolation and messaging to simplify and harden the connectivity between containers Learn about numerous issues encountered at scale and their workarounds, from the kernel up to code versioning Automate the most important parts of your infrastructure with continuous integration In Detail Deploying Docker into production is considered to be one of the major pain points in developing large-scale infrastructures, and the documentation available online leaves a lot to be desired. With this book, you will learn everything you wanted to know to effectively scale your deployments globally and build a resilient, scalable, and containerized cloud platform for your own use. The book starts by introducing you to the containerization ecosystem with some concrete and easy-to-digest examples; after that, you will delve into examples of launching multiple instances of the same container. From there, you will cover orchestration, multi-node setups, volumes, and almost every relevant component of this new approach to deploying services. Using intertwined approaches, the book will cover battle-tested tooling, or issues likely to be encountered in real-world scenarios, in detail. You will also learn about the other supporting components required for a true PaaS deployment and discover common options to tie the whole infrastructure together. At the end of the book, you learn to build a small, but functional, PaaS (to appreciate the power of the containerized service approach) and continue to explore real-world approaches to implementing even larger global-scale services. Style and approach This in-depth learning guide shows you how to deploy your applications in production using Docker (from the basic steps to advanced concepts) and how to overcome challenges in Docker-based infrastructures. The book also covers practical use-cases in real-world examples, and provides tips and tricks on the various topics.

**Deploying Rails** Jun 05 2023 Today's modern Rails applications have lots of moving parts. Make sure your next production deployment goes smoothly with this hands-on book, which guides you through the entire production process.

You'll set up scripts to install and configure all the software your servers need, including your application code. Once you're in production, you'll learn how to set up systems to monitor your application's health, gather metrics so you can stop problems before they start, and fix things when they go wrong. Deploying Rails takes you on an expertly guided tour of the current best practices in Rails deployment and management. You'll find in-depth explanations on effectively running a Rails app by leveraging popular open source tools such as Puppet, Capistrano, and Vagrant. Then you'll go beyond deployment and learn how to use Ganglia and Nagios to monitor your application's health and gather metrics so you can head off problems before they happen. You'll start out by building your own virtual environment by writing scripts to provision a production server with Vagrant and Puppet. Then you'll leverage the popular Rails deployment tool Capistrano to deploy an application into this infrastructure. Once the app is live, you'll monitor your application's health with Nagios, and configure Ganglia to collect system metrics. Finally, you'll see how to keep your data backed up, recover data when things go wrong, tame your log files, and use Puppet to automate everything along the way. Whether you're a Rails developer who wants a better understanding of the needs of a production Rails system, if you're a system administrator who wants to manage a Rails application, or if you're bridging the gap between development and operations, this book will be your roadmap to successful production deployment and maintenance, whether your application has ten users or ten million users. **What You Need:** The exercises and examples are most suited to a computer running some Unix variant, such as Mac OS X or Linux. But a Windows machine running Linux in a VirtualBox virtual machine is also sufficient. We'll show you how to set up a local virtual machine for your deployments; you won't need a dedicated server to hone your deployment skills. We expect you to have a basic familiarity with the Ruby programming language, the Ruby on Rails framework, and the Unix command line.

**OpenShift for Developers** Sep 08 2023 Ready to build cloud native applications? Get a hands-on introduction to daily life as a developer crafting code on OpenShift, the open source container application platform from Red Hat. Creating and packaging your apps for deployment on modern distributed systems can be daunting. Too often, adding infrastructure value can complicate development. With this practical guide, you'll learn how to build, deploy, and manage a multitiered application on OpenShift. Authors Joshua Wood and Brian Tannous demonstrate how OpenShift speeds application development. With the Kubernetes container orchestrator at its core, OpenShift simplifies and automates the way you build, ship, and run code. You'll learn how to use OpenShift and the Quarkus Java framework to develop and deploy apps using proven enterprise technologies and practices that you can apply to code in any language. Learn the development cycles for building and deploying on OpenShift, and the tools that drive them. Use OpenShift to build, deploy, and manage the ongoing lifecycle of an n-tier application. Create a continuous integration and deployment pipeline to turn your source code changes into production rollouts. Automate scaling decisions with metrics and trigger lifecycle events with webhooks.

**Deployment Fundamentals, Vol. 6** Aug 27 2022 If your job is deploying Windows 10, this book is for you. In this book, you will find practical guidance based on our many years of real-world experience deploying Windows around the world. Deployment Fundamentals, Volume 6, provides you with detailed step-by-step instructions, as well as decision-making guidance and explanations that provide answers on the Whys and Hows around Windows 10 OS deployment using Microsoft Deployment Toolkit (MDT) 2013 Update 2. The book also includes many real-world notes and troubleshooting tips and tricks. To get you going as quickly as possible, the book's sample scripts contain a fully automated build of the entire environment, the hydration kit. That includes a fully configured Active Directory environment, including DNS, DHCP, WSUS, PXE, DFS-R Replication, SQL Express, and more. With this book, you will learn how to: Install and configure MDT 2013 Update 2 for production deployments - Build the supporting infrastructure - Use the script repository included with this book in your own environment - Create production-ready reference images for Windows 10 - Build a real-world deployment solution for Windows 10 - Add and deploy applications - Perform real-world driver management - Apply advanced configurations for CustomSettings.ini and deployment automation - Extend MDT using application wrappers, userexit scripts, and PowerShell - Prestage deployment settings using the MDT databases - Perform advanced configuration using web services - Deploy Office 2016, including the Click-to-Run Office 365 version.

**Release It!** Jun 17 2024 A single dramatic software failure can cost a company millions of dollars - but can be avoided with simple changes to design and architecture. This new edition of the best-selling industry standard shows you how to create systems that run longer, with fewer failures, and recover better when bad things happen. New coverage includes DevOps, microservices, and cloud-native architecture. Stability antipatterns have grown to include systemic problems in large-scale systems. This is a must-have pragmatic guide to engineering for production systems. If you're a software developer, and you don't want to get alerts every night for the rest of your life, help is here. With a combination of case studies about huge losses - lost revenue, lost reputation, lost time, lost opportunity - and practical, down-to-earth advice that was all gained through painful experience, this book helps you avoid the pitfalls that cost companies millions of dollars in downtime and reputation. Eighty percent of project life-cycle cost is in production, yet few books address this topic. This updated edition deals with the production of today's systems - larger, more complex, and heavily virtualized - and includes information on chaos engineering, the discipline of applying randomness and deliberate stress to reveal systematic problems. Build systems that survive the real world, avoid downtime, implement zero-downtime upgrades and continuous delivery, and make cloud-native applications resilient. Examine ways to architect, design, and build software - particularly distributed systems - that stands up to the typhoon winds of a flash mob, a Slashdotting, or a link on Reddit. Take a hard look at software that failed the test and find ways to make sure your software survives. To skip the pain and get the experience...get this book.

**Continuous Delivery in Java** Jan 12 2024 Continuous delivery adds enormous value to the business and the entire software delivery lifecycle, but adopting this practice means mastering new skills typically outside of a developer's comfort zone. In this practical book, Daniel Bryant and Abraham Marín-Pérez provide guidance to help experienced Java developers master skills such as architectural design, automated quality assurance, and application packaging and deployment on a variety of platforms. Not only will you learn how to create a comprehensive build pipeline for continually delivering effective software, but you'll also explore how Java application architecture and deployment platforms have affected the way we rapidly and safely deliver new software to production environments. Get advice for beginning or completing your migration to continuous delivery. Design architecture to enable the continuous delivery of Java applications. Build application artifacts including fat JARs, virtual machine images, and operating system container (Docker) images. Use continuous integration tooling like Jenkins, PMD, and find-sec-bugs to automate code quality checks. Create a comprehensive build pipeline and design software to separate the deploy and release processes. Explore why functional and system quality attribute testing is vital from development to delivery. Learn how to effectively build and test applications locally and observe your system while it runs in production.

**Practical Machine Learning with AWS** Apr 10 2021 Successfully build, tune, deploy, and productionize any machine learning model, and know how to automate the process from data processing to deployment. This book is divided into three parts. Part I introduces basic cloud concepts and terminologies related to AWS services such as S3, EC2, Identity Access Management, Roles, Load Balancer, and CloudFormation. It also covers cloud security topics such as AWS Compliance and artifacts, and the AWS Shield and CloudWatch monitoring service built for developers and DevOps engineers. Part II covers machine learning in AWS using SageMaker, which gives developers and data scientists the ability to build, train, and deploy machine learning models. Part III explores other AWS services such as Amazon Comprehend (a natural language processing service that uses machine learning to find insights and relationships in text), Amazon Forecast (helps you deliver accurate forecasts), and Amazon Textract. By the end of the book, you will understand the machine learning pipeline and how to execute any machine learning model using AWS. The book will also help you prepare for the AWS Certified Machine Learning—Specialty certification exam. **What You Will Learn** Be familiar with the different machine learning services offered by AWS. Understand S3, EC2, Identity Access Management, and CloudFormation. Understand SageMaker, Amazon Comprehend, and Amazon Forecast. Execute live projects: from the pre-processing phase to deployment on AWS. **Who This Book Is For** Machine learning engineers who want to learn AWS machine learning services, and acquire an AWS machine learning specialty certification.

**Cloud Foundry: The Definitive Guide** Jun 24 2022 How can Cloud Foundry help you develop and deploy business-critical applications and tasks with velocity? This practical guide demonstrates how this open source, cloud-native application platform not only significantly reduces the develop-to-deploy cycle time, but also raises the value line for application operators by changing the way applications and supporting services are deployed and run. Learn how Cloud Foundry can help you improve your product velocity by handling many of the essential tasks required to run applications in production. Author Duncan Winn shows DevOps and operations teams how to configure and run Cloud Foundry at scale. You'll examine Cloud Foundry's technical concepts—including how various platform components interrelate—and learn how to choose your underlying infrastructure, define the networking architecture, and establish resiliency requirements. This book covers: Cloud-native concepts that make the app build, test, deploy, and scale faster. How to deploy Cloud Foundry and the BOSH release engineering toolchain. Concepts and components

of Cloud Foundry's runtime architecture Cloud Foundry's routing mechanisms and capabilities The platform's approach to container tooling and orchestration BOSH concepts, deployments, components, and commands Basic tools and techniques for debugging the platform Recent and soon-to-emerge features of Cloud Foundry

**Production-Ready Microservices** Dec 11 2023 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

**Production-Ready Applied Deep Learning** Sep 27 2022 Supercharge your skills for developing powerful deep learning models and distributing them at scale efficiently using cloud services Key Features Understand how to execute a deep learning project effectively using various tools available Learn how to develop PyTorch and TensorFlow models at scale using Amazon Web Services Explore effective solutions to various difficulties that arise from model deployment Book Description Machine learning engineers, deep learning specialists, and data engineers encounter various problems when moving deep learning models to a production environment. The main objective of this book is to close the gap between theory and applications by providing a thorough explanation of how to transform various models for deployment and efficiently distribute them with a full understanding of the alternatives. First, you will learn how to construct complex deep learning models in PyTorch and TensorFlow. Next, you will acquire the knowledge you need to transform your models from one framework to the other and learn how to tailor them for specific requirements that deployment environments introduce. The book also provides concrete implementations and associated methodologies that will help you apply the knowledge you gain right away. You will get hands-on experience with commonly used deep learning frameworks and popular cloud services designed for data analytics at scale. Additionally, you will get to grips with the authors' collective knowledge of deploying hundreds of AI-based services at a large scale. By the end of this book, you will have understood how to convert a model developed for proof of concept into a production-ready application optimized for a particular production setting. What you will learn Understand how to develop a deep learning model using PyTorch and TensorFlow Convert a proof-of-concept model into a production-ready application Discover how to set up a deep learning pipeline in an efficient way using AWS Explore different ways to compress a model for various deployment requirements Develop Android and iOS applications that run deep learning on mobile devices Monitor a system with a deep learning model in production Choose the right system architecture for developing and deploying a model Who this book is for Machine learning engineers, deep learning specialists, and data scientists will find this book helpful in closing the gap between the theory and application with detailed examples. Beginner-level knowledge in machine learning or software engineering will help you grasp the concepts covered in this book easily.

**Deploy Machine Learning Models to Production** Feb 13 2024 Build and deploy machine learning and deep learning models in production with end-to-end examples. This book begins with a focus on the machine learning model deployment process and its related challenges. Next, it covers the process of building and deploying machine learning models using different web frameworks such as Flask and Streamlit. A chapter on Docker follows and covers how to package and containerize machine learning models. The book also illustrates how to build and train machine learning and deep learning models at scale using Kubernetes. The book is a good starting point for people who want to move to the next level of machine learning by taking pre-built models and deploying them into production. It also offers guidance to those who want to move beyond Jupyter notebooks to training models at scale on cloud environments. All the code presented in the book is available in the form of Python scripts for you to try the examples and extend them in interesting ways. You will: Build, train, and deploy machine learning models at scale using Kubernetes Containerize any kind of machine learning model and run it on any platform using Docker Deploy machine learning and deep learning models using Flask and Streamlit frameworks.

**Production Ready OpenStack - Recipes for Successful Environments** Apr 03 2023 Over 90 practical and highly applicable recipes to successfully deploy various OpenStack configurations in production About This Book Get a deep understanding of OpenStack's internal structure and services Learn real-world examples on how to build and configure various production grade use cases for each of OpenStack's services Use a step-by-step approach to install and configure OpenStack's services to provide Compute, Storage, and Networking as a services for cloud workloads Who This Book Is For If you have a basic understanding of Linux and Cloud computing and want to learn about configurations that OpenStack supports, this is the book for you. Knowledge of virtualization and managing Linux environments is expected. Prior knowledge or experience of OpenStack is not required, although beneficial. What You Will Learn Plan an installation of OpenStack with a basic configuration Deploy OpenStack in a highly available configuration Configure Keystone Identity services with multiple types of identity backends Configure Glance Image Store with File, NFS, Swift, or Ceph image backends and use local image caching Design Cinder to use a single storage provider such as LVM, Ceph, and NFS backends, or to use multiple storage backends simultaneously Manage and configure the OpenStack networking backend Configure OpenStack's compute hypervisor and the instance scheduling mechanism Build and customize the OpenStack dashboard In Detail OpenStack is the most popular open source cloud platform used by organizations building internal private clouds and by public cloud providers. OpenStack is designed in a fully distributed architecture to provide Infrastructure as a Service, allowing us to maintain a massively scalable cloud infrastructure. OpenStack is developed by a vibrant community of open source developers who come from the largest software companies in the world. The book provides a comprehensive and practical guide to the multiple uses cases and configurations that OpenStack supports. This book simplifies the learning process by guiding you through how to install OpenStack in a single controller configuration. The book goes deeper into deploying OpenStack in a highly available configuration. You'll then configure Keystone Identity Services using LDAP, Active Directory, or the MySQL identity provider and configure a caching layer and SSL. After that, you will configure storage back-end providers for Glance and Cinder, which will include Ceph, NFS, Swift, and local storage. Then you will configure the Neutron networking service with provider network VLANs, and tenant network VXLAN and GRE. Also, you will configure Nova's Hypervisor with KVM, and QEMU emulation, and you will configure Nova's scheduler filters and weights. Finally, you will configure Horizon to use Apache HTTPD and SSL, and you will customize the dashboard's appearance. Style and approach This book consists of clear, concise instructions coupled with practical and applicable recipes that will enable you to use and implement the latest features of OpenStack.

**Kubernetes in Production Best Practices** Feb 01 2023 Design, build, and operate scalable and reliable Kubernetes infrastructure for production Key FeaturesImplement industry best practices to build and manage production-grade Kubernetes infrastructureLearn how to architect scalable Kubernetes clusters, harden container security, and fine-tune resource managementUnderstand, manage, and operate complex business workloads confidentlyBook Description Although out-of-the-box solutions can help you to get a cluster up and running quickly, running a Kubernetes cluster that is optimized for production workloads is a challenge, especially for users with basic or intermediate knowledge. With detailed coverage of cloud industry standards and best practices for achieving scalability, availability, operational excellence, and cost optimization, this Kubernetes book is a blueprint for managing applications and services in production. You'll discover the most common way to deploy and operate Kubernetes clusters, which is to use a public cloud-managed service from AWS, Azure, or Google Cloud Platform (GCP). This book explores Amazon Elastic Kubernetes Service (Amazon EKS), the AWS-managed version of Kubernetes, for working through practical exercises. As you get to grips with implementation details specific to AWS and EKS, you'll understand the design concepts, implementation best practices, and configuration applicable to other cloud-managed services. Throughout the book, you'll also discover standard and cloud-agnostic tools, such as Terraform and Ansible, for provisioning and configuring infrastructure. By the end of this book, you'll be able to leverage Kubernetes to operate and manage your production environments confidently. What you will learnExplore different infrastructure architectures for Kubernetes deploymentImplement optimal open source and commercial storage management solutionsApply best practices for provisioning and configuring Kubernetes clusters, including infrastructure as code (IaC) and configuration as code (CAC)Configure the cluster networking plugin and core networking components to get the best out of themSecure your Kubernetes environment using the latest tools and best practicesDeploy core observability stacks, such as monitoring and logging, to fine-tune your infrastructureWho this book is for This book is for cloud infrastructure experts, DevOps engineers, site reliability engineers, and engineering managers looking to design and operate

Kubernetes infrastructure for production. Basic knowledge of Kubernetes, Terraform, Ansible, Linux, and AWS is needed to get the most out of this book.

**Microservice Patterns and Best Practices** May 12 2021 Explore the concepts and tools you need to discover the world of microservices with various design patterns Key Features Get to grips with the microservice architecture and build enterprise-ready microservice applications Learn design patterns and the best practices while building a microservice application Obtain hands-on techniques and tools to create high-performing microservices resilient to possible fails Book Description Microservices are a hot trend in the development world right now. Many enterprises have adopted this approach to achieve agility and the continuous delivery of applications to gain a competitive advantage. This book will take you through different design patterns at different stages of the microservice application development along with their best practices. Microservice Patterns and Best Practices starts with the learning of microservices key concepts and showing how to make the right choices while designing microservices. You will then move onto internal microservices application patterns, such as caching strategy, asynchronism, CQRS and event sourcing, circuit breaker, and bulkheads. As you progress, you'll learn the design patterns of microservices. The book will guide you on where to use the perfect design pattern at the application development stage and how to break monolithic application into microservices. You will also be taken through the best practices and patterns involved while testing, securing, and deploying your microservice application. At the end of the book, you will easily be able to create interoperable microservices, which are testable and prepared for optimum performance. What you will learn How to break monolithic application into microservices Implement caching strategies, CQRS and event sourcing, and circuit breaker patterns Incorporate different microservice design patterns, such as shared data, aggregator, proxy, and chained Utilize consolidate testing patterns such as integration, signature, and monkey tests Secure microservices with JWT, API gateway, and single sign on Deploy microservices with continuous integration or delivery, Blue-Green deployment Who this book is for This book is for architects and senior developers who would like implement microservice design patterns in their enterprise application development. The book assumes some prior programming knowledge.

**Deep Learning for Coders with fastai and PyTorch** Nov 17 2021 Deep learning is often viewed as the exclusive domain of math PhDs and big tech companies. But as this hands-on guide demonstrates, programmers comfortable with Python can achieve impressive results in deep learning with little math background, small amounts of data, and minimal code. How? With fastai, the first library to provide a consistent interface to the most frequently used deep learning applications. Authors Jeremy Howard and Sylvain Gugger, the creators of fastai, show you how to train a model on a wide range of tasks using fastai and PyTorch. You'll also dive progressively further into deep learning theory to gain a complete understanding of the algorithms behind the scenes. Train models in computer vision, natural language processing, tabular data, and collaborative filtering Learn the latest deep learning techniques that matter most in practice Improve accuracy, speed, and reliability by understanding how deep learning models work Discover how to turn your models into web applications Implement deep learning algorithms from scratch Consider the ethical implications of your work Gain insight from the foreword by PyTorch cofounder, Soumith Chintala

**Real-Time Phoenix** Aug 07 2023 Give users the real-time experience they expect, by using Elixir and Phoenix Channels to build applications that instantly react to changes and reflect the application's true state. Learn how Elixir and Phoenix make it easy and enjoyable to create real-time applications that scale to a large number of users. Apply system design and development best practices to create applications that are easy to maintain. Gain confidence by learning how to break your applications before your users do. Deploy applications with minimized resource use and maximized performance. Real-time applications come with real challenges - persistent connections, multi-server deployment, and strict performance requirements are just a few. Don't try to solve these challenges by yourself - use a framework that handles them for you. Elixir and Phoenix Channels provide a solid foundation on which to build stable and scalable real-time applications. Build applications that thrive for years to come with the best-practices found in this book. Understand the magic of real-time communication by inspecting the WebSocket protocol in action. Avoid performance pitfalls early in the development lifecycle with a catalog of common problems and their solutions. Leverage GenStage to build a data pipeline that improves scalability. Break your application before your users do and confidently deploy them. Build a real-world project using solid application design and testing practices that help make future changes a breeze. Create distributed apps that can scale to many users with tools like Phoenix Tracker. Deploy and monitor your application with confidence and reduce outages. Deliver an exceptional real-time experience to your users, with easy maintenance, reduced operational costs, and maximized performance, using Elixir and Phoenix Channels. What You Need: You'll need Elixir 1.9+ and Erlang/OTP 22+ installed on a Mac OS X, Linux, or Windows machine.

**Learning Serverless** Jul 14 2021 Whether your company is considering serverless computing or has already made the decision to adopt this model, this practical book is for you. Author Jason Katzer shows early- and mid-career developers what's required to build and ship maintainable and scalable services using this model. With this book, you'll learn how to build a modern production system in the cloud, viewed through the lens of serverless computing. You'll discover how serverless can free you from the tedious task of setting up and maintaining systems in production. You'll also explore new ways to level up your career and design, develop, and deploy with confidence. In three parts, this book includes: The Path to Production: Examine the ins and outs of distributed systems, microservices, interfaces, and serverless architecture and patterns The Tools: Dive into monitoring, observability and alerting, logging, pipelines, automation, and deployment Concepts: Learn how to design security and privacy, how to manage quality through testing and staging, and how to plan for failure

**Programming AWS Lambda** Nov 10 2023 Serverless revolutionizes the way organizations build and deploy software. With this hands-on guide, Java engineers will learn how to use their experience in the new world of serverless computing. You'll discover how this cloud computing execution model can drastically decrease the complexity in developing and operating applications while reducing costs and time to market. Engineering leaders John Chapin and Mike Roberts guide you through the process of developing these applications using AWS Lambda, Amazon's event-driven, serverless computing platform. You'll learn how to prepare the development environment, program Lambda functions, and deploy and operate your serverless software. The chapters include exercises to help you through each aspect of the process. Get an introduction to serverless, functions as a service, and AWS Lambda Learn how to deploy working Lambda functions to the cloud Program Lambda functions and learn how the Lambda platform integrates with other AWS services Build and package Java-based Lambda code and dependencies Create serverless applications by building a serverless API and data pipeline Test your serverless applications using automated techniques Apply advanced techniques to build production-ready applications Understand both the gotchas and new opportunities of serverless architecture

**Building Machine Learning Pipelines** Dec 31 2022 Companies are spending billions on machine learning projects, but it's money wasted if the models can't be deployed effectively. In this practical guide, Hannes Hapke and Catherine Nelson walk you through the steps of automating a machine learning pipeline using the TensorFlow ecosystem. You'll learn the techniques and tools that will cut deployment time from days to minutes, so that you can focus on developing new models rather than maintaining legacy systems. Data scientists, machine learning engineers, and DevOps engineers will discover how to go beyond model development to successfully productize their data science projects, while managers will better understand the role they play in helping to accelerate these projects. Understand the steps to build a machine learning pipeline Build your pipeline using components from TensorFlow Extended Orchestrate your machine learning pipeline with Apache Beam, Apache Airflow, and Kubeflow Pipelines Work with data using TensorFlow Data Validation and TensorFlow Transform Analyze a model in detail using TensorFlow Model Analysis Examine fairness and bias in your model performance Deploy models with TensorFlow Serving or TensorFlow Lite for mobile devices Learn privacy-preserving machine learning techniques

**React 17 Design Patterns and Best Practices** May 24 2022 Build scalable, maintainable, and powerful React web apps with design patterns and insightful best practices Key Features Make the most of design patterns in React - including render props and controlled and uncontrolled inputs Master React Hooks with the help of this updated third edition Work through examples that can be used to create reusable code and extensible designs Book Description Filled with useful React patterns that you can use in your projects straight away, this book will help you save time and build better web applications with ease. React 17 Design Patterns and Best Practices is a hands-on guide for those who want to take their coding skills to a new level. You'll spend most of your time working your way through the principles of writing maintainable and clean code, but you'll also gain a deeper insight into the inner workings of React. As you progress through the chapters, you'll learn how to build components that are reusable across the application, how to structure applications, and create forms that actually work. Then you'll build on your knowledge by exploring how to style React components and optimize them to make applications faster and more responsive. Once you've mastered the rest, you'll learn how to write tests effectively and how to contribute to React and its ecosystem. By the end of this book, you'll be able to avoid the process of trial and error and developmental headaches. Instead, you'll be able to use your new skills to efficiently build and deploy real-world React web applications

you can be proud of. What you will learn Get to grips with the techniques of styling and optimizing React components Create components using the new React Hooks Use server-side rendering to make applications load faster Get up to speed with the new React Suspense technique and using GraphQL in your projects Write a comprehensive set of tests to create robust and maintainable code Build high-performing applications by optimizing components Who this book is for This book is for web developers who want to understand React better and apply it to real-life app development. You'll need an intermediate-level experience with React and JavaScript before you get started.

- [Earrings By Judith Viorst](#)
- [California Mathematics Grade 7 Practice Workbook Answers](#)
- [Economic And Financial Decisions Under Risk Exercise Solution](#)
- [Solution Manual Discrete Mathematics And Its Applications 6th Edition](#)
- [Tony Gaddis Java Lab Manual Answers 7th](#)
- [1999 Saturn SL2 Owners Manual](#)
- [1984 Study Guide Answers](#)
- [Ks2 English Targeted Question Grammar Punctuation Spelling Year 5 Cgp Ks2 English](#)
- [Homeland And Other Stories Barbara Kingsolver](#)
- [Solution Manual For Applied Mathematical Programming Bradley](#)
- [Vw Engine Diagram](#)
- [Physics And Everyday Thinking Answer Key](#)
- [Iahcsmm 7th Edition Workbook](#)
- [Plumbing Level 2 Trainee Guide](#)
- [Lannon Technical Communication 12th Edition](#)
- [Introduction To Econometrics Empirical Exercise Solutions](#)
- [Anthropology What Does It Mean To Be Human By Robert H Lavenda And Emily A Schultz Oxford University Press Second Edition](#)
- [Algebra 1 Mcgraw Hill Answers](#)
- [Holt Spanish 2 Assessment Program Answers](#)
- [Adolescence Santrock 15th Edition](#)
- [Oes Worthy Matron Handbook Pdf](#)
- [Creative Writing Four Genres In Brief](#)
- [Criminal Justice An Introduction An Introduction To Crime And The Criminal Justice System](#)
- [High Voltage Engineering Naidu Solution Manual](#)
- [Mitchell 1993 Ford Taurus Sho Repair Manual](#)
- [Stripping Asjiah I](#)
- [How To Escape Your Prison Workbook Answers Pdf](#)
- [Quinox El Angel Oscuro 1 Exilio](#)
- [Pearson Myaccountinglab Answers](#)
- [Ah Bach Math Answers Knowing All Angles](#)
- [Solutions Manual Algorithms Robert Sedgewick 4th Edition](#)
- [Algebra Structure And Method 1 Teacher Edition Online](#)
- [Glencoe Spanish 1 Answer Key](#)
- [Medical Microbiology 6th Edition](#)
- [Repair A Word Document Pdf](#)
- [Prentice Hall Realidades 2 Practice Workbook Answers Key](#)
- [Grammar For Writing Workbook](#)
- [Program Evaluation Test Bank And Solution Manual You](#)
- [Papers On Bullying In Schools](#)
- [1993 Nissan D21 Repair Manual](#)
- [Envision Math Workbook Grade 4 Printable](#)
- [Chloes Kitchen 125 Easy Delicious Recipes For Making The Food You Love Vegan Way Chloe Coscarelli](#)
- [Cleveland Clinic Pbds Study Guide](#)
- [American Government Chapter 4 Federalism](#)
- [Chapter 11 Vocabulary Review Answers](#)
- [Mymathlab Answer Key Elementary Algebra](#)
- [The White Giraffe Questions And Answers](#)
- [Wisconsin Drivers License Template](#)



- [Anatomy And Physiology Coloring Workbook Answer Key Chapter 5](#)
- [Codependent No More Printable](#)