

Web Scalability For Startup Engineers

How do today’s most successful tech companies—Amazon, Google, Facebook, Netflix, Tesla—design, develop, and deploy the products that have earned the love of literally billions of people around the world? Perhaps surprisingly, they do it very differently than the vast majority of tech companies. In INSPIRED, technology product management thought leader Marty Cagan provides readers with a master class in how to structure and staff a vibrant and successful product organization, and how to discover and deliver technology products that your customers will love—and that will work for your business. With sections on assembling the right people and skillsets, discovering the right product, embracing an effective yet lightweight process, and creating a strong product culture, readers can take the information they learn and immediately leverage it within their own organizations—dramatically improving their own product efforts. Whether you’re an early stage startup working to get to product/market fit, or a growth-stage company working to scale your product organization, or a large, long-established company trying to regain your ability to consistently deliver new value for your customers, INSPIRED will take you and your product organization to a new level of customer engagement, consistent innovation, and business success. Filled with the author’s own personal stories—and profiles of some of today’s most-successful product managers and technology-powered product companies, including Adobe, Apple, BBC, Google, Microsoft, and Netflix—INSPIRED will show you how to turn up the dial of your own product efforts, creating technology products your customers love. The first edition of INSPIRED, published ten years ago, established itself as the primary reference for technology product managers, and can be found on the shelves of nearly every successful technology product company worldwide. This thoroughly updated second edition shares the same objective of being the most valuable resource for technology product managers, yet it is completely new—sharing the latest practices and techniques of today’s most-successful tech product companies, and the men and women behind every great product.

Apply design principles to your classes, preparing them for reuse. You will use package design principles to create packages that are just right in terms of cohesion and coupling, and are user- and maintainer-friendly at the same time. The first part of this book walks you through the five SOLID principles that will help you improve the design of your classes. The second part introduces you to the best practices of package design, and covers both package cohesion principles and package coupling principles. Cohesion principles show you which classes should be put together in a package, when to split packages, and if a combination of classes may be considered a "package" in the first place. Package coupling principles help you choose the right dependencies and prevent wrong directions in the dependency graph of your packages. What You'll LearnApply the SOLID principles of class designDetermine if classes belong in the same packageKnow whether it is safe for packages to depend on each other Who This Book Is For Software developers with a broad range of experience in the field, who are looking for ways to reuse,share, and distribute their code

If you need to build a scalable, fault tolerant system with requirements for high availability, discover why the Erlang/OTP platform stands out for the breadth, depth, and consistency of its features. This hands-on guide demonstrates how to use the Erlang programming language and its OTP framework of reusable libraries, tools, and design principles to develop complex commercial-grade systems that simply cannot fail. In the first part of the book, you'll learn how to design and implement process behaviors and supervision trees with Erlang/OTP, and bundle them into standalone nodes. The second part addresses reliability, scalability, and high availability in your overall system design. If you're familiar with Erlang, this book will help you understand the design choices and trade-offs necessary to keep your system running. Explore OTP's building blocks: the Erlang language, tools and libraries collection, and its abstract principles and design rules Dive into the fundamentals of OTP reusable frameworks: the Erlang process structures OTP uses for behaviors Understand how OTP behaviors support client-server structures, finite state machine patterns, event handling, and runtime/code integration Write your own behaviors and special processes Use OTP's tools, techniques, and architectures to handle deployment, monitoring, and operations

An introduction to Oracle's use of open source technology explores the applications of nearly one hundred open source tools, including such common applications as Linux, Perl, and Apache, as well as Oracle-specific ones such as Orasof, Orac, OracleTool, and OraSnap, and provides information on how to get them, their benefits, and more. Original. (Intermediate)

Why a book about logs? That’s easy: the humble log is an abstraction that lies at the heart of many systems, from NoSQL databases to cryptocurrencies. Even though most engineers don’t think much about them, this short book shows you why logs are worthy of your attention. Based on his popular blog posts, LinkedIn principal engineer Jay Kreps shows you how logs work in distributed systems, and then delivers practical applications of these concepts in a variety of common uses—data integration, enterprise architecture, real-time stream processing, data system design, and abstract computing models. Go ahead and take the plunge with logs; you’re going love them. Learn how logs are used for programmatic access in databases and distributed systems Discover solutions to the huge data integration problem when more data of more varieties meet more systems Understand why logs are at the heart of real-time stream processing Learn the role of a log in the internals of online data systems Explore how Jay Kreps applies these ideas to his own work on data infrastructure systems at LinkedIn

Software Performance and Scalability

Designing Web APIs

Site Reliability Engineering

Be More Than Just a Code Monkey

Scalability Rules

Designing Fine-Grained Systems

The Art of Scalability

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

Want your web site to display more quickly? This book presents 14 specific rules that will cut 25% to 50% off response time when users request a page. Author Steve Souders, in his job as Chief Performance Yahoo!, collected these best practices while optimizing some of the most-visited pages on the Web. Even sites that had already been highly optimized, such as Yahoo! Search and the Yahoo! Front Page, were able to benefit from these surprisingly simple performance guidelines. The rules in High Performance Web Sites explain how you can optimize the performance of the Ajax, CSS, JavaScript, Flash, and images that you've already built into your site -- adjustments that are critical for any rich web application. Other sources of information pay a lot of attention to tuning web servers, databases, and hardware, but the bulk of display time is taken up on the browser side and by the communication between server and browser. High Performance Web Sites covers every aspect of that process. Each performance rule is supported by specific examples, and code snippets are available on the book's companion web site. The rules include how to: Make Fewer HTTP Requests Use a Content Delivery Network Add an Expires Header Gzip Components Put Stylesheets at the Top Put Scripts at the Bottom Avoid CSS Expressions Make JavaScript and CSS External Reduce DNS Lookups Minify JavaScript Avoid Redirects Remove Duplicates Scripts Configure ETags Make Ajax Cacheable If you're building pages for high traffic destinations and want to optimize the experience of users visiting your site, this book is indispensable. "If everyone would implement just 20% of Steve's guidelines, the Web would be dramatically better place. Between this book and Steve's YSlow extension, there's really no excuse for having a sluggish web site anymore." -Joe Hewitt, Developer of Firebug debugger and Mozilla's DOM Inspector "Steve Souders has done a fantastic job of distilling a massive, semi-arcanic art down to a set of concise, actionable, pragmatic engineering steps that will change the world of web performance." -Eric Lawrence, Developer of the Fiddler Web Debugger, Microsoft Corporation

This invaluable roadmap for startup engineers reveals how to successfully handle web application scalability challenges to meet increasing product and traffic demands. Web Scalability for Startup Engineers shows engineers working at startups and small companies how to plan and implement a comprehensive scalability strategy. It presents broad and holistic view of infrastructure and architecture of a scalable web application. Successful startups often face the challenge of scalability, and the core concepts driving a scalable architecture are language and platform agnostic. The book covers scalability of HTTP-based systems (websites, REST APIs, SaaS, and mobile application backends), starting with a high-level perspective before taking a deep dive into common challenges and issues. This approach builds a holistic view of the problem, helping you see the big picture, and then introduces different technologies and best practices for solving the problem at hand. The book is enriched with the author's real-world experience and expert advice, saving you precious time and effort by learning from others' mistakes and successes. Language-agnostic approach addresses universally challenging concepts in Web development/scalability—does not require knowledge of a particular language Fills the gap for engineers in startups and smaller companies who have limited means for getting to the next level in terms of accomplishing scalability Strategies presented help to decrease time to market and increase the efficiency of web applications

Writing code is the easy part of your work as a software developer. This practical book lets you explore the other 90%—everything from requirements discovery and rapid prototyping to business analysis and designing for maintainability. Instead of providing neatly packaged advice from on high, author Gregory Brown presents detailed examples of the many problems developers encounter, including the thought process it takes to solve them. He does this in an unusual and entertaining fashion by making you the main character in a series of chapter-length stories. As these stories progress, the examples become more complex, and your responsibilities increase. Together, these stories take you on a journey that will make you question and refine the way you think about, and work on, software projects. Steps in this unique journey include: Using prototypes to explore project ideas Spotting hidden dependencies in incremental changes Identifying the pain points of service integrations Developing a rigorous approach towards problem-solving Designing software from the bottom up Data modeling in an imperfect world Gradual process improvement as an antidote for over-commitment The future of software development

Describing how infrastructure and software architecture work together to support a scalable environment; this book focuses on core concepts and best practices; rather than on individual languages; platforms; or technologies; and explains step by step how scalable systems work and how to solve common challenges. --

Database Reliability Engineering

I Heart Logs

Scalable Web Architecture, Processes, and Organizations for the Modern Enterprise

The Startup Owner's Manual

Creating Reusable Software Components

Best Practices for Improving Embedded Systems Development

Modern Front-end Architecture

Summary Node.js in Action, Second Edition is a thoroughly revised book based on the best-selling first edition. It starts at square one and guides you through all the features, techniques, and concepts you'll need to build production-quality Node applications. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology You already know JavaScript. The trick to mastering Node.js is learning how to build applications that fully exploit its powerful asynchronous event handling and non-blocking I/O features. The Node server radically simplifies event-driven real-time apps like chat, games, and live data analytics, and with its incredibly rich ecosystem of modules, tools, and libraries, it's hard to beat! About the Book Based on the bestselling first edition, Node.js in Action, Second Edition is a completely new book. Packed with practical examples, it teaches you how to create high-performance web servers using JavaScript and Node. You'll master key design concepts such as asynchronous programming, state management, and event-driven programming. And you'll learn to put together MVC servers using Express and Connect, design web APIs, and set up the perfect production environment to build, lint, and test. What's Inside Mastering non-blocking I/O The Node event loop Testing and deploying Web application templating About the Reader Written for web developers with intermediate JavaScript skills. About the Authors The Second Edition author team includes Node masters Alex Young, Bradley Meck, Mike Cantelon, and Tim Oxley, along with original authors Marc Harter, T.J. Holowaychuk, and Nathan Rajlich. Table of contents PART 1 - WELCOME TO NODE Welcome to Node.js Node programming fundamentals What is a Node web application? PART 2 - WEB DEVELOPMENT WITH NODE Front-end build systems Server-side frameworks Connect and Express in depth Web application templating Storing application data Testing Node applications Deploying Node applications and maintaining uptime PART 3 - BEYOND WEB DEVELOPMENT Writing command-line applications Conquering the desktop with Electron

Using a web API to provide services to application developers is one of the more satisfying endeavors that software engineers undertake. But building a popular API with a thriving developer ecosystem is also one of the most challenging. With this practical guide, developers, architects, and tech leads will learn how to navigate complex decisions for designing, scaling, marketing, and evolving interoperable APIs. Authors Brenda Jin, Saurabh Sahni, and Amir Shevat explain API design theory and provide hands-on exercises for building your web API and managing its operation in production. You'll also learn how to build and maintain a following of app developers. This book includes expert advice, worksheets, checklists, and case studies from companies including Slack, Stripe, Facebook, Microsoft, Cloudinary, Oracle, and GitHub. Get an overview of request-response and event-driven API design paradigms Learn best practices for designing an API that meets the needs of your users Use a template to create an API design process Scale your web API to support a growing number of API calls and use cases Regularly adapt the API to reflect changes to your product or business Provide developer resources that include API documentation, samples, and tools

Why care about hardware/firmware interaction? These interfaces are critical, a solid hardware design married with adaptive firmware can access all the capabilities of an application and overcome limitations caused by poor communication. For the first time, a book has come along that will help hardware engineers and firmware engineers work together to mitigate or eliminate problems that occur when hardware and firmware are not optimally compatible. Solving these issues will save time and money, getting products to market sooner to create more revenue. The principles and best practices presented in this book will prove to be a valuable resource for both hardware and firmware engineers. Topics include register layout, interrupts, timing and performance, aborts, and errors. Real world cases studies will help to solidify the principles and best practices with an aim towards cleaner designs, shorter schedules, and better implementation! Reduce product development delays with the best practices in this book Concepts apply to ASICs, ASSPs, SoCs, and FPGAs Real-world examples and case studies highlight the good and bad of design processes

Annotation Over the past 10 years, distributed systems have become more fine-grained. From the large multi-million line long monolithic applications, we are now seeing the benefits of smaller self-contained services. Rather than heavy-weight, hard to change Service Oriented Architectures, we are now seeing systems consisting of collaborating microservices. Easier to change, deploy, and if required retire, organizations which are in the right position to take advantage of them are yielding significant benefits. This book takes an holistic view of the things you need to be cognizant of in order to pull this off. It covers just enough understanding of technology, architecture, operations and organization to show you how to move towards finer-grained systems.

Fully updated! Fifty Powerful, Easy-to-Use Rules for Supporting Hyper Growth “Whether you’re taking on a role as a technology leader in a new company or you simply want to make great technology decisions, Scalability Rules will be the go-to resource on your bookshelf.”

—Chad Dickerson, CTO, Etsy Scalability Rules, Second Edition, is the easy-to-use scalability primer and reference for every architect, developer, network/software engineer, web professional, and manager. Authors Martin L. Abbott and Michael T. Fisher have helped scale hundreds of high-growth companies and thousands of systems. Drawing on their immense experience, they present 50 up-to-the-minute technical best practices for supporting hyper growth practically anywhere. Fully updated to reflect new technical trends and experiences, this edition is even easier to read, understand, and apply. Abbott and Fisher have also added powerful “stories behind the rules”: actual experiences and case studies from CTOs and technology executives at Etsy, NASDAQ, Salesforce, Shutterfly, Chegg, Warby Parker, Twitter, and other scalability pioneers. Architects will find powerful technology-agnostic insights for creating and evaluating designs. Developers will discover specific techniques for handling everything from databases to state. Managers will get invaluable help in setting goals, making decisions, and interacting with technical teams. Whatever your role, you'll find practical risk/benefit guidance for setting priorities, translating plans into action, and gaining maximum scalability at minimum cost. You'll learn how to Simplify architectures and avoid “over-engineering” Design scale into your solution, so you can scale on a just-in-time basis Make the most of cloning and replication Separate functionality and split data sets Scale out, not up Get more out of databases without compromising scalability Eliminate unnecessary redirects and redundant double-checking Use caches and CDNs more aggressively, without unacceptable complexity Design for fault tolerance, graceful failure, and easy rollback Emphasize statelessness, and efficiently handle state when you must Effectively utilize asynchronous communication Learn from your own mistakes and others' high-profile failures Prioritize your actions to get the biggest “bang for the buck”

Microservice Patterns and Best Practices

Building Mobile Apps at Scale

Designing for Scalability with Erlang/OTP

Web Scalability for Startup Engineers

Programming Beyond Practices

Blitzscaling

Tips & Techniques for Scaling Your Web Application

Praise from the Reviewers: “The practicality of the subject in a real-world situation distinguishes this book from others available on the market.” —Professor Behrouz Far, University of Calgary “This book could replace the computer organization texts now in use that every CS and CpE student must take. . . . It is much needed, well written, and thoughtful.” —Professor Larry Bernstein, Stevens Institute of Technology A distinctive, educational text onsoftware performance and scalability This is the first book to take a quantitative approach to the subject of software performance and scalability. It brings together three unique perspectives to demonstrate how your products can be optimized and tuned for the best possible performance and scalability: The Basics—introduces the computer hardware and software architectures that predetermine the performance and scalability of a software product as well as the principles of measuring the performance and scalability of a software product Queuing Theory—helps you learn the performance laws and queuing models for interpreting the underlying physics behind software performance and scalability, supplemented with ready-to-apply techniques for improving the performance and scalability of a software system API Profiling—shows you how to design more efficient algorithms and achieve optimized performance and scalability, aided by adopting an API profiling framework (perfBasic) built on the concept of a performance map for drilling down performance root causes at the API level Software Performance and Scalability gives you a specialized skill set that will enable you to design and build performance into your products with immediate, measurable improvements. Complemented with real-world case studies, it is an indispensable resource for software developers, quality and performance assurance engineers, architects, and managers. It is anideal text for university courses related to computer and software performance evaluation and can also be used to supplement a course in computer organization or in queuing theory for upper-division and graduate computer science students.

In-depth examination of concepts and principles of Web application development Completely revised and updated, this popular book returns with coverage on a range of new technologies. Authored by a highly respected duo, this edition provides an in-depth examination of the core concepts and general principles of Web application development. Packed with examples featuring specific technologies, this book is divided into three sections: HTTP protocol as a foundation for Web applications, markup languages (HTML, XML, and CSS), and survey of emerging technologies. After a detailed introduction to the history of Web applications, coverage segues to core Internet protocols, Web browsers, Web application development, trends and directions, and more. Includes new coverage on technologies such as application primers, Ruby on Rails, SOAP, XPath, P3P, and more Explores the fundamentals of HTTP and its evolution Looks at HTML and its roots as well as XML languages and applications Reviews the basic operation of Web Servers, their functionality, configuration, and security Discusses how to process flow in Web browsers and looks at active browser pages Addresses the trends and various directions that the future of Web application frameworks may be headed This book is essential reading for anyone who needs to design or debug complex systems, and it makes it easier to learn the new application programming interfaces that arise in a rapidly changing Internet environment.

Learn how to build front-end applications that can help you ship applications faster with fewer defects. Many software projects fail because they are not planned well, or lack organization. Applying strategies from other industries can help you create better software. This book explores the “mise en place” technique from cooking and reveals how you can apply it to the art of creating software. In many professional kitchens, the work of preparing a meal is divided and performed in workstations to help create the whole. You'll review the philosophy behind this and see how to leverage this approach in your code. You'll also learn to visualize your applications as a series of components and build parts of your application in isolation - just like a professional chef crafting each part of the meal. The book describes to how to structure your code base for reuse, and how to communicate the code's intent to other developers. You'll develop your components in isolation and test these building blocks for quality at a granular level. Then compose these components as building blocks in increasingly complicated features. Finally, you'll apply some strategies not directly related to code to ensure maximum quality and efficiency. With Modern Front-end Architecture, developers of all levels will learn strategies that they and their teams can leverage to be more productive. What You'll Learn Structure an application as a series of components Build a component library that others in an organization can leverage Ensure quality and accessibility at a component level rather than a page or app level Test code in a way that gives the maximum amount of confidence while providing an excellent developer experience Automate repeatable tasks Who This Book Is For Front end developers looking to maximize their code for reuse, quality, and shipping features quickly. Experienced developers will pick up new techniques that they can apply to their code base. Less experienced developers will be able to start applications off on the right foot.

Publisher's Note: Products purchased from Third Party sellers are not guaranteed by the publisher for quality, authenticity, or access to any online entitlements included with the product. Design and build scalable web applications quickly This is an invaluable roadmap for meeting the rapid demand to deliver scalable applications in a startup environment. With a focus on core concepts and best practices rather than on individual languages, platforms, or technologies, Web Scalability for Startup Engineers describes how infrastructure and software architecture work together to support a scalable environment. You'll learn, step by step, how scalable systems work and how to solve common challenges. Helpful diagrams are included throughout, and real-world examples illustrate the concepts presented. Even if you have limited time and resources, you can successfully develop and deliver

robust, scalable web applications with help from this practical guide. Learn the key principles of good software design required for scalable systems Build the front-end layer to sustain the highest levels of concurrency and request rates Design and develop web services, including REST-ful APIs Enable a horizontally scalable data layer Implement caching best practices Leverage asynchronous processing, messaging, and event-driven architecture Structure, index, and store data for optimized search Explore other aspects of scalability, such as automation, project management, and agile teams

Are you working on a codebase where cost overruns, death marches, and heroic fights with legacy code monsters are the norm? Battle these adversaries with novel ways to identify and prioritize technical debt, based on behavioral data from how developers work with code. And that's just for starters. Because good code involves social design, as well as technical design, you can find surprising dependencies between people and code to resolve coordination bottlenecks among teams. Best of all, the techniques build on behavioral data that you already have: your version-control system. Join the fight for better code! Use statistics and data science to uncover both problematic code and the behavioral patterns of the developers who build your software. This combination gives you insights you can't get from the code alone. Use these insights to prioritize refactoring needs, measure their effect, find implicit dependencies between different modules, and automatically create knowledge maps of your system based on actual code contributions. In a radical, much-needed change from common practice, guide organizational decisions with objective data by measuring how well your development teams align with the software architecture. Discover a comprehensive set of practical analysis techniques based on version-control data, where each point is illustrated with a case study from a real-world codebase. Because the techniques are language neutral, you can apply them to your own code no matter what programming language you use. Guide organizational decisions with objective data by measuring how well your development teams align with the software architecture. Apply research findings from social psychology to software development, ensuring you get the tools you need to coach your organization towards better code. If you're an experienced programmer, software architect, or technical manager, you'll get a new perspective that will change how you work with code. What You Need: You don't have to install anything to follow along in the book. TThe case studies in the book use well-known open source projects hosted on GitHub. You'll use CodeScene, a free software analysis tool for open source projects, for the case studies. We also discuss alternative tooling options where they exist.

The Big Ideas Behind Reliable, Scalable, and Maintainable Systems
Essential Knowledge for Front-End Engineers
Best Practices for Designing High Volume Websites
Fix Technical Debt with Behavioral Code Analysis
Building APIs That Developers Love

How to Create Tech Products Customers Love

The system design interview is considered to be the most complex and most difficult technical job interview by many. Those questions are intimidating, but don't worry. It's just that nobody has taken the time to prepare you systematically of diagrams and use lots of examples. You'll learn step-by-step, one question at a time.Don't miss out.What's inside?- An insider's take on what interviewers really look for and why.- A 4-step framework for solving any system design interview interview questions with detailed solutions.- 188 diagrams to visually explain how different systems work.

Foreword by Bill Gates LinkedIn cofounder, legendary investor, and host of the award-winning Masters of Scale podcast reveals the secret to starting and scaling massively valuable companies. What entrepreneur or founder doesn't aspire to Airbnb? Yet those who actually manage to do so are exceedingly rare. So what separates the startups that get disrupted and disappear from the ones who grow to become global giants? The secret is blitzscaling: a set of techniques for scaling competitors out of the water. The objective of Blitzscaling is not to go from zero to one, but from one to one billion –as quickly as possible. When growing at a breakneck pace, getting to next level requires very different strategies from a book inspired by their popular class at Stanford Business School, Hoffman and Yeh reveal how to navigate the necessary shifts and weather the unique challenges that arise at each stage of a company's life cycle, such as: how to design a relentless growth: strategies for hiring and managing; how the role of the founder and company culture must evolve as the business matures, and more. Whether your business has ten employees or ten thousand, Blitzscaling is the essential speed is the only competitive advantage that matters.

More than 100,000 entrepreneurs rely on this book for detailed, step-by-step instructions on building successful, scalable, profitable startups. The National Science Foundation pays hundreds of startup teams each year to follow the process Stanford, Berkeley, Columbia and more than 100 other leading universities worldwide. Why? The Startup Owner's Manual guides you, step-by-step, as you put the Customer Development process to work. This method was created by renowned Blank, co-creator with Eric Ries of the "Lean Startup" movement and tested and refined by him for more than a decade. This 608-page how-to guide includes over 100 charts, graphs, and diagrams, plus 77 valuable checklists that guide you to profitability. It will help you: • Avoid the 9 deadly sins that destroy startups' chances for success • Use the Customer Development method to bring your business idea to life • Incorporate the Business Model Canvas as the organizing principle for customers and determine how to "get, keep and grow" customers profitably • Compute how you'll drive your startup to repeatable, scalable profits. The Startup Owner's Manual was originally published by K&S Ranch Publishing Inc. and is now a design, and content are the same as the prior release and should not be considered a new or updated product.

The Comprehensive, Proven Approach to IT Scalability-Updated with New Strategies, Technologies, and Case Studies In The Art of Scalability, Second Edition, leading scalability consultants Martin L. Abbott and Michael T. Fisher cover everything you scale products and services for any requirement. This extensively revised edition reflects new technologies, strategies, and lessons, as well as new case studies from the authors' pioneering consulting practice, AKF Partners. Writing for tech Abbott and Fisher cover everything that impacts scalability, including architecture, process, people, organization, and technology. Their insights and recommendations reflect more than thirty years of experience at companies ranging from e-commerce to You'll find updated strategies for structuring organizations to maximize agility and scalability, as well as new insights into the cloud (IaaS/PaaS) transition, NoSQL, DevOps, business metrics, and more. Using this guide's tools and advice, you can overcome obstacles to scalability-and achieve unprecedented IT and business performance. Coverage includes • Why scalability problems start with organizations and people, not technology, and what to do about it • Actionable lessons from real success stories and leading the agile, scalable organization • Scaling processes for hyper-growth environments • Architecting scalability: proprietary models for clarifying needs and making choices-including 15 key success principles • Emerging technologies for planning, cloud evolution, and customer-aligned monitoring • Measuring availability, capacity, load, and performance

A guide to developing Web sites using scalable applications.

Scalability Patterns

Optimize Your Front-end Development with Components, Storybook, and Mise en Place Philosophy

Principles of Package Design

Production-Ready Microservices

Hello, Startup

Hardware/Firmware Interface Design

Oracle and Open Source

Cloud native infrastructure is more than servers, network, and storage in the cloud—it is as much about operational hygiene as it is about elasticity and scalability. In this book, you'll learn practices, patterns, and requirements for creating infrastructure that meets your needs, capable of managing the full life cycle of cloud native applications. Justin Garrison and Kris Nova reveal hard-earned lessons on architecting infrastructure from companies such as Google, Amazon, and Netflix. They draw inspiration from projects adopted by the Cloud Native Computing Foundation (CNCF), and provide examples of patterns seen in existing tools such as Kubernetes. With this book, you will: Understand why cloud native infrastructure is necessary to effectively run cloud native applications Use guidelines to decide when—and if—your business should adopt cloud native practices Learn patterns for deploying and managing infrastructure and applications Design tests to prove that your infrastructure works as intended, even in a variety of edge cases Learn how to secure infrastructure with policy as code Great software architects aren't born. They are a product of decades of building real-life solutions and relentless learning. They become really good at their trade closer to the retirement age. But most startups are fostered by young entrepreneurs who dare to try but lack the experience. They also lack the \$\$ to hire a silver-haired architect to join their team from day one. Left to their own faculties, the entrepreneurs and their engineering teams quickly get on the path of learning from their own mistakes. Eventually, they discover this is the most expensive way of learning. Over time they get better, and some become the true masters of the craft - but way too late to make a difference for their early-day projects.This book is meant to break the vicious circle. It isn't a textbook, at least not in the traditional sense. It is a business-centric practical guide to software architecture, intended for software engineers, technology executives, students of computer science, and tech-savvy entrepreneurs who want to de-risk their entrepreneurial endeavors or to fast-track their careers in software engineering. The recipes in this book are highly practical, battle-tested, and current for building mid- to large-scale systems in 2019.

Data is at the center of many challenges in system design today. Difficult issues need to be figured out, such as scalability, consistency, reliability, efficiency, and maintainability. In addition, we have an overwhelming variety of tools, including relational databases, NoSQL datastores, stream or batch processors, and message brokers. What are the right choices for your application? How do you make sense of all these buzzwords? In this practical and comprehensive guide, author Martin Kleppmann helps you navigate this diverse landscape by examining the pros and cons of various technologies for processing and storing data. Software keeps changing, but the fundamental principles remain the same. With this book, software engineers and architects will learn how to apply those ideas in practice, and how to make full use of data in modern applications. Peer under the hood of the systems you already use, and learn how to use and operate them more effectively Make informed decisions by identifying the strengths and weaknesses of different tools Navigate the trade-offs around consistency, scalability, fault tolerance, and complexity Understand the distributed systems research upon which modern databases are built Peek behind the scenes of major online services, and learn from their architectures

API development is becoming increasingly common for server-side developers thanks to the rise of front-end JavaScript frameworks, iPhone applications, and API-centric architectures. It might seem like grabbing stuff from a data source and shoving it out as JSON would be easy, but surviving changes in business logic, database schema updates, new features, or deprecated endpoints can be a nightmare. After finding many of the existing resources for API development to be lacking, Phil learned a lot of things the hard way through years of trial and error. This book aims to condense that experience, taking examples and explanations further than the trivial apples and pears nonsense tutorials often provide. By passing on some best practices and general good advice you can hit the ground running with API development, combined with some horror stories and how they were overcome/avoided/averted. This book will discuss the theory of designing and building APIs in any language or framework, with this theory applied in PHP-based examples.

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.

39 Engineering Challenges

Design Patterns for Cloud Native Applications

Become an Awesome Software Architect

System Design Interview - An Insider's Guide

Principles for Scaling Web Sites

Everyone and Their Dog Wants an API, So You Should Probably Learn How to Build Them

Event Data, Stream Processing, and Data Integration

With the immense cost savings and scalability the cloud provides, the rationale for building cloud native applications using APIs, data, events, and streams in both greenfield and brownfield development. You'll learn how to incrementally design, develop, and deploy large and effective cloud native applications that you can manage and maintain at scale with minimal cost, time, and effort. Authors Kasun Indrasiri and Srisikandarajah Suhothayan highlight use cases that effectively demonstrate the challenges you might encounter at each step. Learn the fundamentals of cloud native applications Explore key cloud native communication, connectivity, and composition patterns Learn decentralized data management techniques Use event-driven architecture to build distributed and scalable cloud native applications Explore the most commonly used patterns for API management and consumption Examine some of the tools and technologies you'll need for building cloud native systems

Summary Machine Learning Systems: Designs that scale is an example-rich guide that teaches you how to implement reactive design solutions in your machine learning systems to make them as reliable as a well-built web app. Foreword by Sean Owen, Director of Data Science, Cloudera Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology If you're building machine learning models to be used on a small scale, you don't need this book. But if you're a developer building a production-grade ML application that needs quick response times, reliability, and good user experience, this is the book for you. It collects principles and practices of machine learning systems that are dramatically easier to run and maintain, and that are reliably better for users. About the Book Machine Learning Systems: Designs that scale teaches you to design and implement production-ready ML systems. You'll learn the principles of reactive design as you build pipelines with Spark, create highly scalable services with Akka, and use powerful machine learning libraries like MLlib on massive datasets. The examples use the Scala language, but the same ideas and tools work in Java, as well. What's Inside Working with Spark, MLlib, and Akka Reactive design patterns Monitoring and maintaining a large-scale system Futures, actors, and supervision About the Reader Readers need intermediate skills in Java or Scala. No prior machine learning experience is assumed. About the Author Jeff Smith builds powerful machine learning systems. For the past decade, he has been working on building data science applications, teams, and companies as part of various teams in New York, San Francisco, and Hong Kong. He blogs (https://medium.com/@jeffksmithjr), tweets (@jeffksmithjr), and speaks (www.jeffsmith.tech/speaking) about various aspects of building real-world machine learning systems. Table of Contents PART 1 - FUNDAMENTALS OF REACTIVE MACHINE LEARNING Learning reactive machine learning Using reactive tools PART 2 - BUILDING A REACTIVE MACHINE LEARNING SYSTEM Collecting data Generating features Learning models Evaluating models Publishing models Responding PART 3 - OPERATING A MACHINE LEARNING SYSTEM Delivering Evolving intelligence

50 Powerful, Easy-to-Use Rules for Supporting Hypergrowth in Any Environment Scalability Rules is the easy-to-use scalability primer and reference for every architect, developer, web professional, and manager. Authors Martin L. Abbott and Michael T. Fisher have helped scale more than 200 hypergrowth Internet sites through their consulting practice. Now, drawing on their unsurpassed experience, they present 50 clear, proven scalability rules—and practical guidance for applying them. Abbott and Fisher transform scalability from a “black art” to a set of realistic, technology-agnostic best practices for supporting hypergrowth in nearly any environment, including both frontend and backend systems. For architects, they offer powerful new insights for creating and evaluating designs. For developers, they share specific techniques for handling everything from databases to state. For managers, they provide invaluable help in goal-setting, decision-making, and interacting with technical teams. Whatever your role, you'll find practical risk/benefit guidance for setting priorities—and getting maximum “bang for the buck.” • Simplifying architectures and avoiding “over-engineering” • Scaling via cloning, replication, separating functionality, and splitting data sets • Scaling out, not up • Getting more out of databases without compromising scalability • Avoiding unnecessary redirects and redundant double-checking • Using caches and content delivery networks more aggressively, without introducing unacceptable complexity • Designing for fault tolerance, graceful failure, and easy rollback • Striving for statelessness when you can; efficiently handling state when you must • Effectively utilizing asynchronous communication • Learning quickly from mistakes, and much more

A Comprehensive, Proven Approach to IT Scalability from Two Veteran Software, Technology, and Business Executives In The Art of Scalability, AKF Partners cofounders Martin L. Abbott and Michael T. Fisher cover everything IT and business leaders must know to build technology infrastructures that can scale smoothly to meet any business requirement. Drawing on their unparalleled experience managing some of the world's highest-transaction-volume Web sites, the authors provide detailed models and best-practice approaches available in no other book. Unlike previous books on scalability, The Art of Scalability doesn't limit its coverage to technology. Writing for both technical and nontechnical decision-makers, this book covers everything that impacts scalability, including architecture, processes, people, and organizations. Throughout, the authors address a broad spectrum of real-world challenges, from performance testing to IT governance. Using their tools and guidance, organizations can systematically overcome obstacles to scalability and achieve unprecedented levels of technical and business performance. Coverage includes Staffing the scalable organization: essential organizational, management, and leadership skills for technical leaders Building processes for scale: process lessons from hyper-growth companies, from technical issue resolution to crisis management Making better “build versus buy” decisions Architecting scalable solutions: powerful proprietary models for identifying scalability needs and choosing the best approaches to meet them Optimizing performance through caching, application and database splitting, and asynchronous design Scalability techniques for emerging technologies, including clouds and grids Planning for rapid data growth and new data centers Evolving monitoring strategies to tightly align with customer requirements

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

97 Things Every Cloud Engineer Should Know

Principles, Protocols and Practices

A Programmer's Guide to Building Products, Technologies, and Teams

Build APIs You Won't Hate

Node.js in Action

Building Scalable Web Sites

Building Standardized Systems Across an Engineering Organization

If you create, manage, operate, or configure systems running in the cloud, you're a cloud engineer—even if you work as a system administrator, software developer, data scientist, or site reliability engineer. With this book, professionals from around the world provide valuable insight into today's cloud engineering role. These concise articles explore the entire cloud computing experience, including fundamentals, architecture, and migration. You'll delve into security and compliance, operations and reliability, and software development. And examine networking, organizational culture, and more. You're sure to find 1, 2, or 97 things that inspire you to dig deeper and expand your own career. “Three Keys to Making the Right Multicloud Decisions,” Brendan O’Leary “Serverless Bad Practices,” Manases Jesus Galindo Bello “Failing a Cloud Migration,” Lee Atchison “Treat Your Cloud Environment as If It Were On Premises,” Ilyana Garry “What Is Toil, and Why Are SREs Obsessed with It?,” Zachary Nickens “Lean QA: The QA Evolving in the DevOps World,” Theresa Neate “How Economies of Scale Work in the Cloud,” Jon Moore “The Cloud Is Not About the Cloud,” Ken Corless “Data Gravity: The Importance of Data Management in the Cloud,” Geoff Hughes “Even in the Cloud, the Network Is the Foundation,” David Murray “Cloud Engineering Is About Culture, Not Containers,” Holly Cummins

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

In this book, the CEO of Cazton, Inc. and internationally-acclaimed speaker, Chander Dhall, demonstrates current website design scalability patterns and takes a pragmatic approach to explaining their pros and cons to show you how to select the appropriate pattern for your site. He then tests the patterns by deliberately forcing them to fail and exposing potential flaws before discussing how to design the optimal pattern to match your scale requirements. The author explains the use of polyglot programming and how to match the right patterns to your business needs. He also details several No-SQL patterns and explains the fundamentals of different paradigms of No-SQL by showing complementary strategies of using them along with relational databases to achieve the best results. He also teaches how to make the scalability pattern work with a real-world microservices pattern. With the proliferation of countless electronic devices and the ever growing number of Internet users, the scalability of websites has become an increasingly important challenge. Scalability, even though highly coveted, may not be so easy to achieve. Think that you can't attain responsiveness along with scalability? Chander Dhall will demonstrate that, in fact, they go hand in hand. What You'll Learn Architect and develop applications so that

they are easy to scale. Learn different scaling and partitioning options and the combinations. Learn techniques to speed up responsiveness. Deep dive into caching, column-family databases, document databases, search engines and RDBMS. Learn scalability and responsiveness concepts that are usually ignored. Effectively balance scalability, performance, responsiveness, and availability while minimizing downtime. Who This Book Is For Executives (CXOs), software architects, developers, and IT Pros

While there is a lot of appreciation for backend and distributed systems challenges, there tends to be less empathy for why mobile development is hard when done at scale. This book collects challenges engineers face when building iOS and Android apps at scale, and common ways to tackle these. By scale, we mean having numbers of users in the millions and being built by large engineering teams. For mobile engineers, this book is a blueprint for modern app engineering approaches. For non-mobile engineers and managers, it is a resource with which to build empathy and appreciation for the complexity of world-class mobile engineering. The book covers iOS and Android mobile app challenges on these dimensions: Challenges due to the unique nature of mobile applications compared to the web, and to the backend. App complexity challenges. How do you deal with increasingly complicated navigation patterns? What about non-deterministic event combinations? How do you localize across several languages, and how do you scale your automated and manual tests? Challenges due to large engineering teams. The larger the mobile team, the more challenging it becomes to ensure a consistent architecture. If your company builds multiple apps, how do you balance not rewriting everything from scratch while moving at a fast pace, over waiting on "centralized" teams? Cross-platform approaches. The tooling to build mobile apps keeps changing. New languages, frameworks, and approaches that all promise to address the pain points of mobile engineering keep appearing. But which approach should you choose? Flutter, React Native, Cordova? Native apps? Reuse business logic written in Kotlin, C#, C++ or other languages? What engineering approaches do "world-class" mobile engineering teams choose in non-functional aspects like code quality, compliance, privacy, compliance, or with experimentation, performance, or app size?

Web Scalability for Startup Engineers McGraw Hill Professional

The Lightning-Fast Path to Building Massively Valuable Companies

Book 1: Foundation 2019

High Performance Web Sites

A Quantitative Approach

Machine Learning Systems

How Google Runs Production Systems

Explore patterns like CORS and event sourcing to create scalable, maintainable, and testable microservices

The infrastructure-as-code revolution in IT is also affecting database administration. With this practical book, developers, system administrators, and junior to mid-level DBAs will learn how the modern practice of site reliability engineering applies to the craft of database architecture and operations. Authors Laine Campbell and Charity Majors provide a framework for professionals looking to join the ranks of today's database reliability engineers (DBRE). You'll begin by exploring core operational concepts that DBREs need to master. Then you'll examine a wide range of database persistence options, including how to implement key technologies to provide resilient, scalable, and performant data storage and retrieval. With a firm foundation in database reliability engineering, you'll be ready to dive into the architecture and operations of any modern database. This book covers: Service-level requirements and risk management Building and evolving an architecture for operational visibility Infrastructure engineering and infrastructure management How to facilitate the release management process Data storage, indexing, and replication Identifying datastore characteristics and best use cases Datastore architectural components and data-driven architectures

Implement Robust, Fault-Tolerant Systems

Building Microservices

Cloud Native Infrastructure

50 Principles for Scaling Web Sites

Designing and Operating Resilient Database Systems

Patterns for Scalable Infrastructure and Applications in a Dynamic Environment

INSPIRED