

# Web Scalability For Startup Engineers

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 need to know to smoothly 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 technical and nontechnical decision-makers, 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 eBay to Visa, and Salesforce.com to Apple. 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 systematically clear away 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 successes and failures • Staffing, structuring, 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 and challenges: data cost, datacenter planning, cloud evolution, and customer-

aligned monitoring • Measuring availability, capacity, load, and performance

Looks at the application design process, describing how to create user-friendly applications.

Today, software engineers need to know not only how to program effectively but also how to develop proper engineering practices to make their codebase sustainable and healthy. This book emphasizes this difference between programming and software engineering. How can software engineers manage a living codebase that evolves and responds to changing requirements and demands over the length of its life? Based on their experience at Google, software engineers Titus Winters and Hyrum Wright, along with technical writer Tom Manshreck, present a candid and insightful look at how some of the world's leading practitioners construct and maintain software. This book covers Google's unique engineering culture, processes, and tools and how these aspects contribute to the effectiveness of an engineering organization. You'll explore three fundamental principles that software organizations should keep in mind when designing, architecting, writing, and maintaining code: How time affects the sustainability of software and how to make your code resilient over time How scale affects the viability of software practices within an engineering organization What trade-offs a typical engineer needs to make when evaluating design and development decisions

TensorFlow is one of the most popular machine learning frameworks in Python. With this book, you will improve your knowledge of some of the latest TensorFlow features and will be able to perform supervised and unsupervised machine learning and also train neural networks.

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 build the next Amazon, Facebook, or 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 up at a dizzying pace that blows 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 those that got you to where you are today. In 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 business models for igniting and sustaining 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 playbook for winning in a world where speed is the only competitive advantage that matters.

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 MLib 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, MLib, 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](http://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

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.

Web Scalability for Startup Engineers McGraw Hill Professional

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 to 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

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

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 Summary MongoDB in Action, Second Edition is a completely revised and updated version. It introduces MongoDB 3.0 and the document-oriented database model. This perfectly paced book gives you both the big picture you'll need as a developer and enough low-level detail to satisfy system engineers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology This document-oriented database was built for high availability, supports rich, dynamic schemas, and lets you easily distribute data across multiple



servers. MongoDB 3.0 is flexible, scalable, and very fast, even with big data loads. About the Book MongoDB in Action, Second Edition is a completely revised and updated version. It introduces MongoDB 3.0 and the document-oriented database model. This perfectly paced book gives you both the big picture you'll need as a developer and enough low-level detail to satisfy system engineers. Lots of examples will help you develop confidence in the crucial area of data modeling. You'll also love the deep explanations of each feature, including replication, auto-sharding, and deployment. What's Inside Indexes, queries, and standard DB operations Aggregation and text searching Map-reduce for custom aggregations and reporting Deploying for scale and high availability Updated for Mongo 3.0 About the Reader Written for developers. No previous MongoDB or NoSQL experience is assumed. About the Authors After working at MongoDB, Kyle Banker is now at a startup. Peter Bakkum is a developer with MongoDB expertise. Shaun Verch has worked on the core server team at MongoDB. A Genentech engineer, Doug Garrett is one of the winners of the MongoDB Innovation Award for Analytics. A software architect, Tim Hawkins has led search engineering at Yahoo Europe. Technical Contributor: Wouter Thielen. Technical Editor: Mihalis Tsoukalos. Table of Contents PART 1 GETTING STARTED A database

for the modern web MongoDB through the JavaScript shell Writing programs using MongoDB  
PART 2 APPLICATION DEVELOPMENT IN MONGODB Document-oriented data Constructing queries Aggregation Updates, atomic operations, and deletes PART 3 MONGODB MASTERY Indexing and query optimization Text search WiredTiger and pluggable storage Replication Scaling your system with sharding Deployment and administration

Have you ever... - Wanted to work at an exciting futuristic company? - Struggled with an interview problem that could have been solved in 15 minutes? - Wished you could study real-world computing problems? If so, you need to read Elements of Programming Interviews (EPI). EPI is your comprehensive guide to interviewing for software development roles. The core of EPI is a collection of over 250 problems with detailed solutions. The problems are representative of interview questions asked at leading software companies. The problems are illustrated with 200 figures, 300 tested programs, and 150 additional variants. The book begins with a summary of the nontechnical aspects of interviewing, such as strategies for a great interview, common mistakes, perspectives from the other side of the table, tips on negotiating the best offer, and a guide to the best ways to use EPI. We also provide a summary of data structures, algorithms, and problem

solving patterns. Coding problems are presented through a series of chapters on basic and advanced data structures, searching, sorting, algorithm design principles, and concurrency. Each chapter starts with a brief introduction, a case study, top tips, and a review of the most important library methods. This is followed by a broad and thought-provoking set of problems. A practical, fun approach to computer science fundamentals, as seen through the lens of common programming interview questions. Jeff Atwood/Co-founder, Stack Overflow and Discourse

**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

America's greatness comes from people working hard to fulfill their dreams. But today that greatness is being undermined by people using the government to steal other people's dreams (and money). Rather than participate and innovate in the marketplace, generating goods and services that benefit society, people are increasingly vying for political advantage to live at the expense of others. Something for Nothing reveals the social and personal threats inherent in this emerging "grabbing

match" culture, juxtaposing free-market virtues against government vices, explaining how the something-for-nothing mentality corrupts the political system, undermines corporate success, and stifles the individual's ability to prosper and contribute long-term to society. More than exposing the dangers, however, Tracy helps readers set a personal and culture-wide agenda for change.

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

Taught at junior level math courses at every university, Linear Algebra is essential for students in almost every technical and analytic discipline. 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

Every day, companies struggle to scale critical applications. As traffic volume and data demands increase, these applications become more complicated and brittle, exposing risks and compromising availability. With the popularity of software as a service, scaling has never been more important. Updated with an expanded focus on modern architecture paradigms such as microservices and cloud computing, this practical guide provides techniques for building systems that can handle huge quantities of traffic, data, and demand—without affecting the quality your customers expect. Architects, managers, and directors in engineering and operations organizations will learn how to build applications at scale that run more smoothly and reliably to meet the needs of customers. Learn how scaling affects the availability of your services, why that matters, and how to improve it

Dive into a modern service-based application architecture that ensures high availability and reduces the effects of service failures Explore the Single Team Owned Service Architecture paradigm (STOSA)—a model for scaling your development organization in tandem with your application Understand, measure, and mitigate risk in your systems Use the cloud to build highly scalable

applications

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.

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.

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

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

Summary Entity Framework Core in Action teaches you how to access and update relational data from .NET applications. Following the crystal-clear explanations, real-world examples, and around 100 diagrams, you'll discover time-saving patterns and best practices for security, performance tuning, and unit testing. Purchase of the print book includes a free eBook in PDF, Kindle,

and ePub formats from Manning Publications. About the Technology There's a mismatch in the way OO programs and relational databases represent data. Entity Framework is an object-relational mapper (ORM) that bridges this gap, making it radically easier to query and write to databases from a .NET application. EF creates a data model that matches the structure of your OO code so you can query and write to your database using standard LINQ commands. It will even automatically generate the model from your database schema. About the Book Using crystal-clear explanations, real-world examples, and around 100 diagrams, Entity Framework Core in Action teaches you how to access and update relational data from .NET applications. You'll start with a clear breakdown of Entity Framework, long with the mental model behind ORM. Then you'll discover time-saving patterns and best practices for security, performance tuning, and even unit testing. As you go, you'll address common data access challenges and learn how to handle them with Entity Framework. What's Inside Querying a relational database with LINQ Using EF Core in business logic Integrating EF with existing C# applications Applying domain-driven design to EF Core Getting the best performance out of EF Core Covers EF Core 2.0 and 2.1 About the Reader For .NET developers with some awareness of how relational databases work. About the Author Jon P Smith is a full-stack developer with special focus on .NET Core and Azure. Table of Contents Part 1 - Getting started Introduction to Entity FrameworkCore Querying the database Changing the database content Using EF Core in business logic Using

EF Core in ASP.NET Core web applications Part 2 - Entity Framework in depth Configuring nonrelational properties Configuring relationships Configuring advanced features and handling concurrency conflicts Going deeper into the DbContext Part 3 - Using Entity Framework Core in real-world applications Useful software patterns for EF Core applications Handling database migrations EF Core performance tuning A worked example of performance tuning Different database types and EF Core services Unit testing EF Core applications Appendix A - A brief introduction to LINQ Appendix B - Early information on EF Core version 2.1

In the race to compete in today's fast-moving markets, large enterprises are busy adopting new technologies for creating new products, processes, and business models. But one obstacle on the road to digital transformation is placing too much emphasis on technology, and not enough on the types of processes technology enables. What if different lines of business could build their own services and applications—and decision-making was distributed rather than centralized? This report explores the concept of a digital business platform as a way of empowering individual business sectors to act on data in real time. Much innovation in a digital enterprise will increasingly happen at the edge, whether it involves business users (from marketers to data scientists) or IoT devices. To facilitate the process, your core IT team can provide these sectors with the digital tools they need to innovate quickly. This report explores: Key cultural and organizational changes for developing business

capabilities through cross-functional product teams A platform for integrating applications, data sources, business partners, clients, mobile apps, social networks, and IoT devices Creating internal API programs for building innovative edge services in low-code or no-code environments Tools including Integration Platform as a Service, Application Platform as a Service, and Integration Software as a Service The challenge of integrating microservices and serverless architectures Event-driven architectures for processing and reacting to events in real time You'll also learn about a complete pervasive integration solution as a core component of a digital business platform to serve every audience in your organization.

This book presents a new theory explaining underdevelopment in the global South and tests whether financial inputs, the government-business-media (GBM) complex and spatiotemporal influences drive human development. Despite the entrance of emerging powers and new forms of aid, trade and investment, international political-economic practices still support well-established systems of capital accumulation, to the detriment of the global South. Global asymmetrical accumulation is maintained by 'affective' (consent-forming hegemonic practices) and 'infrastructural' (uneven economic exchanges) labours and by power networks. The message for developing countries is that 'robust' GBMs can facilitate human development and development is constrained by spatiotemporal limitations. This work theorizes that aid and foreign direct investment should be viewed with caution and that in the global South these

investments should not automatically be assumed to be drivers of development.

This goal of this book is to provide a reliable and easy to understand strategy to approach system design questions. The process and justification of your ideas are the most important things in system design interviews. Thus the combination of right strategy and knowledge is vital to the success of your interview. Some candidates fail because lack of knowledge while some fail because they do not find the right way to approach the problem. This book provides valuable ways to fix both problems. By the time you finish the book, you are exceptionally well-equipped to tackle any system design questions.

About the author Alex is an experienced software engineer and entrepreneur. He enjoys hand-on engineering and the thrill of working on a variety of software products including business applications, web apps and mobile apps. He has worked at Apple and Twitter among other internet companies. While not doing software development, Alex enjoys hiking and gaming. During the job interviews, he learned many things about system design interviews and achieved many successes. But, it is very time consuming to find the effective materials to prepare the interview, so Alex wrote this book offering the best knowledge to ace the design interviews. Alex hopes this book will save you a lot of time, energy to master the system design questions.

TABLE OF CONTENTS  
CHAPTER ONE: SCALE FROM ZERO TO TEN MILLION USERS  
CHAPTER TWO: DESIGN CONSISTENT HASHING  
CHAPTER THREE: DESIGN A KEY-VALUE STORE  
CHAPTER FOUR:

## DESIGN A URL SHORTENER

Outlines a revisionist approach to management while arguing against common perceptions about the inevitability of startup failures, explaining the importance of providing genuinely needed products and services as well as organizing a business that can adapt to continuous customer feedback. 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 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. 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. --

Department of Startup: Why Every Fortune 500 Should Have One aims to help CEOs, presidents, and human resource practitioners on how they can transform their large corporation to thrive in a fast, social media conscious, and unforgiving market a la startup. The ascent of startups on the Fortune 500 ranking, displacing some of the more notable companies, raised not only attention but also questions. Fundamentally, startups are built on a very different organizational culture as compared to a traditional Fortune 500. Could these cultural differences be the reason startups are in the forefront of technology innovation, disrupting industries dominated by more established competitors and thriving in today's volatile business environment?

Department of Startup: Why Every Fortune 500 Should Have One aims to help CEOs, presidents, and human resource

practitioners on how they can transform their large corporation to thrive in a fast, social media conscious, and unforgiving market a la startup.

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"

This is the eBook of the printed book and may not include any media, website access codes, or print supplements that may come packaged with the bound book. Conquer today's Windows 10—from the inside out! Dive into Windows 10—and really put your Windows expertise to work. Focusing on the most powerful and innovative features of Windows 10, this supremely organized reference packs hundreds of timesaving solutions, tips, and workarounds—all fully reflecting the major Windows 10 Anniversary Update. From new Cortana and Microsoft Edge enhancements to the latest security and virtualization features, you'll discover how experts tackle today's essential tasks—and challenge yourself to new levels of mastery. Install, configure, and personalize the newest versions of Windows 10 Understand Microsoft's revamped activation and upgrade processes Discover major Microsoft Edge enhancements, including new support for extensions Use today's improved Cortana services to perform tasks, set reminders, and retrieve information Make the most of the improved ink, voice, touch, and gesture support in Windows 10 Help secure Windows 10 in business with Windows Hello and Azure AD Deploy, use, and manage new Universal Windows Platform (UWP) apps Take advantage of new entertainment options, including Groove Music Pass subscriptions and connections to your Xbox One console Manage files in the cloud with Microsoft OneDrive and

OneDrive for Business Use the improved Windows 10 Mail and Calendar apps and the new Skype app Fine-tune performance and troubleshoot crashes Master high-efficiency tools for managing Windows 10 in the enterprise Leverage advanced Hyper-V features, including Secure Boot, TPMs, nested virtualization, and containers In addition, this book is part of the Current Book Service from Microsoft Press. Books in this program will receive periodic updates to address significant software changes for 12 to 18 months following the original publication date via a free Web Edition. Learn more at <https://www.microsoftpressstore.com/cbs>.

This book provides a complete understanding of the jBPM technology stack. It starts with an introduction to the world of business process management systems, the problem domain addressed by jBPM, explores the main use cases that can be addressed by business process management systems, and illustrates the main design patterns. It takes you through the details of the architecture and available out-of-the-box provisions for customizing, extending, and integrating the features of jBPM to meet the requirements of your application. Moreover, this book will empower you with the knowledge to integrate jBPM with enterprise architecture, debug through the source code of jBPM, and utilize the flexibility provided by a heavily modular system. Finally, it introduces you to the provisions available for a jBPM-based application to put the non-functional characteristics of the system, which are of great importance when we deploy our application in production. The book helps you in putting the knowledge at work by providing you with a lot of ready to use examples, both basic and advanced ones.

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

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 outlined in the book, and it's taught at 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 Silicon Valley startup expert Steve 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 as you drive your company toward 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 startup hypotheses • Identify your 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 available from Wiley. The cover, design, and content are the same as the prior release and should not be considered a new or updated product.

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

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.

Summary Revised and updated for Elixir 1.7, Elixir in Action, Second Edition teaches you how to apply Elixir to practical problems associated with scalability, fault tolerance, and high availability. Along the way, you'll develop an appreciation for, and considerable skill in, a functional and concurrent style of

programming. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology When you're building mission-critical software, fault tolerance matters. The Elixir programming language delivers fast, reliable applications, whether you're building a large-scale distributed system, a set of backend services, or a simple web app. And Elixir's elegant syntax and functional programming mindset make your software easy to write, read, and maintain. About the Book Elixir in Action, Second Edition teaches you how to build production-quality distributed applications using the Elixir programming language. Author Saša Jurišić introduces this powerful language using examples that highlight the benefits of Elixir's functional and concurrent programming. You'll discover how the OTP framework can radically reduce tedious low-level coding tasks. You'll also explore practical approaches to concurrency as you learn to distribute a production system over multiple machines. What's inside Updated for Elixir 1.7 Functional and concurrent programming Introduction to distributed system design Creating deployable releases About the Reader You'll need intermediate skills with client/server applications and a language like Java, C#, or Ruby. No previous experience with Elixir required. About the Author Saša Jurišić is a developer with extensive experience using Elixir and Erlang in complex server-side systems. Table of Contents First steps Building blocks Control flow Data abstractions Concurrency primitives Generic server processes Building a concurrent system Fault-tolerance basics Isolating error effects Beyond GenServer Working with components Building a distributed system Running the system

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



