

Api Management Springer

This book provides readers the “big picture” and a comprehensive survey of the domain of big data processing systems. For the past decade, the Hadoop framework has dominated the world of big data processing, yet recently academia and industry have started to recognize its limitations in several application domains and thus, it is now gradually being replaced by a collection of engines that are dedicated to specific verticals (e.g. structured data, graph data, and streaming data). The book explores this new wave of systems, which it refers to as Big Data 2.0 processing systems. After Chapter 1 presents the general background of the big data phenomena, Chapter 2 provides an overview of various general-purpose big data processing systems that allow their users to develop various big data processing jobs for different application domains. In turn, Chapter 3 examines various systems that have been introduced to support the SQL flavor on top of the Hadoop infrastructure and provide competing and scalable performance in the processing of large-scale structured data. Chapter 4 discusses several systems that have been designed to tackle the problem of large-scale graph processing, while the main focus of Chapter 5 is on several systems that have been designed to provide scalable solutions for processing big data streams, and on other sets of systems that have been introduced to support the development of data pipelines between various types of big data processing jobs and systems. Next, Chapter 6 focuses on covering the emerging frameworks and systems in the domain of scalable machine learning and deep learning processing. Lastly, Chapter 7 shares conclusions and an outlook on future research challenges. This new and considerably enlarged second edition not only contains the completely new chapter 6, but also offers a refreshed content for the state-of-the-art in all domains of big data processing over the last years. Overall, the book offers a valuable reference guide for professional, students, and researchers in the domain of big data processing systems. Further, its comprehensive content will hopefully encourage readers to pursue further research on the subject.

Develop RESTful web services using the Flask micro-framework and integrate them using MySQL. Use Flask to develop, deploy, and manage REST APIs with easy-to-read and understand Python code. Solve your problem from a choice of libraries. Learn to use MySQL as the web services database for your Flask API using SQLAlchemy ORM. Building REST APIs with Flask provides a primer on Flask, RESTful services, and working with pip to set up your virtual environment. The key differences between NoSQL and SQL are covered, and you are taught how to connect MySQL and Flask using SQLAlchemy. Author Kunal Relan presents best practices for creating REST APIs and guides you in structuring your app and testing REST endpoints. He teaches you how to set up authentication and render HTML using views. You learn how to write unit tests for your REST APIs, and understand mocks, assertions, and integration testing. You will know how to document your REST APIs, deploy your Flask application on all of the major cloud platforms, and debug and monitor your Flask application. What You'll Learn Use MySQL to create Flask REST APIs Test REST endpoints Create CRUD endpoints with Flask and MySQL Deploy Flask on all of the major cloud platforms Monitor your Flask application Who This Book Is For Python developers interested in REST API development using Flask and web

developers with basic programming knowledge who want to learn how Python and REST APIs work together. Readers should be familiar with Python (command line, or at least pip) and MySQL.

This textbook explores the theoretical and practical aspects of managing international business operations while also dealing with multi-cultural, multi-national and global issues of managing business expansion beyond the domestic market. A second, revised edition of *Managing Internationally: Succeeding in a Culturally Diverse World*, each chapter contains up-to-date material, in-depth coverage of topics, visual aids (i.e., charts, tables, etc.), and vignettes, making this new edition engaging, visually appealing and easily accessible for students taking International Business Management courses. The contents of this textbook are separated into four parts. Part one offers introductory information on the scope and importance of international business management as well as the social and ethical challenges. Part two covers cultural and behavioral topics. Part three discusses the strategic and operational aspects of international business management. Part four explores human resources and labor relations. To assist students, each chapter starts a preview section which includes an outline of the chapter indicating the important aspects along with a brief description of the major issues. Following the preview is a vignette that encapsulates the crux of the chapter, often presented in an amusing and engaging manner. To further help students focus on key issues, the text includes the list of useful business cases to which students can refer. To assist professors in teaching from this book, ancillary teaching materials such as sample syllabi, slides, tests and answer keys will be available for download.

The Facebook API allows web developers to create Facebook applications and access Facebook data from other applications. Facebook API Developers Guide covers the use and implementation of the Facebook API—what the key features are and how you can access them. You will learn, through practical examples, the main features of the Facebook API including an introduction to the API-specific languages FQL and FBML. These examples are further supported by the introduction of other technologies like language libraries, relational database management systems, and XML. Covers all key features of the Facebook API Explains the API languages FQL and FBML Teaches by example, with useful code and tips you can use in your own applications

Learn to design, implement, measure, and improve DevOps programs that are tailored to your organization. This concise guide assists leaders who are accountable for the rapid development of high-quality software applications. In *DevOps for Digital Leaders*, deep collective experience on both sides of the dev-ops divide informs the global thought leadership and penetrating insights of the authors, all three of whom are cross-portfolio DevOps leaders at CA Technologies. Aruna Ravichandran, Kieran Taylor, and Peter Waterhouse analyze the organizational benefits, costs, freedoms, and constraints of DevOps. They chart the coordinated strategy of organizational change, metrics, lean thinking, and investment that an enterprise must undertake to realize the full potential of DevOps and reach the sweet spot where accelerating code deployments drive increasing customer satisfaction, revenue, and profitability. Digital leaders are charged to bridge the dev-ops disconnect if their organizations are to survive and flourish in a business world increasingly differentiated by the degree to which dynamic application software development harmonizes with operational resilience and reliability. This short book applies the DevOps perspective to the competitive

challenge, faced by every high-performance IT organization today, of integrating and automating open source, cloud, and enterprise tools, processes, and techniques across the software development life cycle from requirements to release. What You Will Learn: Remove dependencies and constraints so that parallel practices can accelerate the development of defect-free software Automate continuous delivery across the software life cycle to eliminate release bottlenecks, manual labor waste, and technical debt accumulation Generate virtualized production-style testing of applications through real-time behavioral analytics Adopt agile practices so operations teams can support developer productivity with automated feedback, streamline infrastructure monitoring, spot and resolve operations issues before they impact production, and improve customer experience Identify the DevOps metrics appropriate to your organization and integrate DevOps with your existing best practices and investment Who This Book Is For: IT leaders in large companies and government agencies who have any level of responsibility for the rapid development of high-quality software applications. The secondary readership is members of development and operations teams, security professionals, and service managers.

This book contains the refereed proceedings of the 12th International Conference on Knowledge Management in Organizations, KMO 2017, held in Beijing, China, in August 2017. The theme of the conference was "Emerging Technology and Knowledge Management in Organizations." The 45 contributions accepted for KMO 2017 were selected from 112 submissions and are organized in topical sections on: Knowledge Management Models and Behaviour Studies; Knowledge Sharing; Knowledge Transfer and Learning; Knowledge and Service Innovation; Knowledge and Organization; Information Systems Research; Value Chain and Supply Chain; Knowledge Re-presentation and Reasoning; Data Mining and Intelligent Science; Big Data Management; Internet of Things and Network.

Maximize the impact of your assets and business services by providing APIs for developers and other users. The journey described in this book starts with identifying business assets. As part of the API team, you then need to identify and define the requirements of traffic management, security, mediation, and orchestration. You also must define metrics for the analytics to measure the success of the overall API program. API documentation and the ease of developer onboarding also determine the success of the APIs. Finally, monetization of these APIs leads to revenue generation for the enterprise. Author De — an expert in building and managing API solutions — provides enterprise architects, designers, and technologists with insight into the world of APIs and the various technical aspects of building and managing an effective API management solution. API Management: Developing and Managing APIs for your Organization: Introduces the basics of APIs and highlights their value Provides an overview of technologies for building an API management solution and defines the requirements, including how to build a RESTful API Offers design principles for building developer-friendly APIs Explains how to secure your APIs Shows how to use API analytics to measure the success of your APIs Demonstrates how to monetize APIs Finally, API Management touches on various technical nuances of creating, distributing, and managing an API. This book will not only help you learn how to design, build, deploy, and manage an API for an enterprise scale, but also generate revenue for your organization. What You'll Learn Discover the API life

cycle Design and develop APIs Implement API security Test your APIs Deploy and monitor your APIs Who This Book Is For Enterprise architects, technology enthusiasts, security architects, and operations specialists.

Prepare for the next wave of challenges in enterprise security. Learn to better protect, monitor, and manage your public and private APIs. Enterprise APIs have become the common way of exposing business functions to the outside world. Exposing functionality is convenient, but of course comes with a risk of exploitation. This book teaches you about TLS Token Binding, User Managed Access (UMA) 2.0, Cross Origin Resource Sharing (CORS), Incremental Authorization, Proof Key for Code Exchange (PKCE), and Token Exchange. Benefit from lessons learned from analyzing multiple attacks that have taken place by exploiting security vulnerabilities in various OAuth 2.0 implementations. Explore root causes, and improve your security practices to mitigate against similar future exploits. Security must be an integral part of any development project. This book shares best practices in designing APIs for rock-solid security. API security has evolved since the first edition of this book, and the growth of standards has been exponential. OAuth 2.0 is the most widely adopted framework that is used as the foundation for standards, and this book shows you how to apply OAuth 2.0 to your own situation in order to secure and protect your enterprise APIs from exploitation and attack. What You Will Learn Securely design, develop, and deploy enterprise APIs Pick security standards and protocols to match business needs Mitigate security exploits by understanding the OAuth 2.0 threat landscape Federate identities to expand business APIs beyond the corporate firewall Protect microservices at the edge by securing their APIs Develop native mobile applications to access APIs securely Integrate applications with SaaS APIs protected with OAuth 2.0 Who This Book Is For Enterprise security architects who are interested in best practices around designing APIs. The book is also for developers who are building enterprise APIs and integrating with internal and external applications.

Business Information Systems: Concepts, Methodologies, Tools and Applications offers a complete view of current business information systems within organizations and the advancements that technology has provided to the business community. This four-volume reference uncovers how technological advancements have revolutionized financial transactions, management infrastructure, and knowledge workers.

You might think more than enough design books exist in the programming world already. In fact, there are so many that it makes sense to ask why you would read yet another. Is there really a need for yet another design book? In fact, there is a greater need than ever before, and Practical API Design: Confessions of a Java Framework Architect fills that need! Teaches you how to write an API that will stand the test of time Written by the designer of the NetBeans API at Sun Technologies Based on best practices, scalability, and API design patterns

Get well-versed with FastAPI features and best practices for testing, monitoring, and deployment to run high-quality and robust data science applications Key Features Cover the concepts of the FastAPI framework, including aspects relating to asynchronous programming, type hinting, and dependency injection Develop efficient RESTful APIs for data science with modern Python Build, test, and deploy high performing data science and machine learning systems with FastAPI Book Description FastAPI is a web

framework for building APIs with Python 3.6 and its later versions based on standard Python-type hints. With this book, you'll be able to create fast and reliable data science API backends using practical examples. This book starts with the basics of the FastAPI framework and associated modern Python programming language concepts. You'll be taken through all the aspects of the framework, including its powerful dependency injection system and how you can use it to communicate with databases, implement authentication and integrate machine learning models. Later, you'll cover best practices relating to testing and deployment to run a high-quality and robust application. You'll also be introduced to the extensive ecosystem of Python data science packages. As you progress, you'll learn how to build data science applications in Python using FastAPI. The book also demonstrates how to develop fast and efficient machine learning prediction backends and test them to achieve the best performance. Finally, you'll see how to implement a real-time face detection system using WebSockets and a web browser as a client. By the end of this FastAPI book, you'll have not only learned how to implement Python in data science projects but also how to maintain and design them to meet high programming standards with the help of FastAPI. What you will learn

- Explore the basics of modern Python and async I/O programming
- Get to grips with basic and advanced concepts of the FastAPI framework
- Implement a FastAPI dependency to efficiently run a machine learning model
- Integrate a simple face detection algorithm in a FastAPI backend
- Integrate common Python data science libraries in a web backend
- Deploy a performant and reliable web backend for a data science application

Who this book is for This Python data science book is for data scientists and software developers interested in gaining knowledge of FastAPI and its ecosystem to build data science applications. Basic knowledge of data science and machine learning concepts and how to apply them in Python is recommended.

This book constitutes the refereed proceedings of the 10th International Conference on Software Business, ICSOB 2019, held in Jyväskylä, Finland, in November 2019. On the occasion of its tenth anniversary the conference theme this year was “The First Decade and Beyond” and focused on the development during the past decade, addressing the future of software-intensive business as well as studies on new and emerging ideas. The 18 full papers and 10 short papers presented together with 3 invited talks, 6 emerging research papers and a tutorial were carefully reviewed and selected from 52 submissions. They are organized in the following topical sections: software ecosystems; management of software products; continual improvement and product development; impacts of digitalization; software business education; software startups and digital business.

This book constitutes the thoroughly refereed proceedings of the 7th International Conference on Data Management Technologies and Applications, DATA 2018, held in Porto, Portugal, in July 2018. The 9 revised full papers were carefully reviewed and selected from 69 submissions. The papers deal with the following topics: databases, big data, data mining, data management, data security, and other aspects of information systems and technology involving advanced applications of data.

Building Your Next Big Thing with Google Cloud Platform shows you how to take advantage of the Google Cloud Platform technologies to build all kinds of cloud-hosted software and services for both public and private consumption. Whether you need a simple virtual server to run your legacy application or you need to architect a sophisticated high-traffic web application, Cloud

Platform provides all the tools and products required to create innovative applications and a robust infrastructure to manage them. Google is known for the scalability, reliability, and efficiency of its various online products, from Google Search to Gmail. And, the results are impressive. Google Search, for example, returns results literally within fractions of second. How is this possible? Google custom-builds both hardware and software, including servers, switches, networks, data centers, the operating system's stack, application frameworks, applications, and APIs. Have you ever imagined what you could build if you were able to tap the same infrastructure that Google uses to create and manage its products? Now you can! Building Your Next Big Thing with Google Cloud Platform shows you how to take advantage of the Google Cloud Platform technologies to build all kinds of cloud-hosted software and services for both public and private consumption. Whether you need a simple virtual server to run your legacy application or you need to architect a sophisticated high-traffic web application, Cloud Platform provides all the tools and products required to create innovative applications and a robust infrastructure to manage them. Using this book as your compass, you can navigate your way through the Google Cloud Platform and turn your ideas into reality. The authors, both Google Developer Experts in Google Cloud Platform, systematically introduce various Cloud Platform products one at a time and discuss their strengths and scenarios where they are a suitable fit. But rather than a manual-like "tell all" approach, the emphasis is on how to Get Things Done so that you get up to speed with Google Cloud Platform as quickly as possible. You will learn how to use the following technologies, among others: Google Compute Engine Google App Engine Google Container Engine Google App Engine Managed VMs Google Cloud SQL Google Cloud Storage Google Cloud Datastore Google BigQuery Google Cloud Dataflow Google Cloud DNS Google Cloud Pub/Sub Google Cloud Endpoints Google Cloud Deployment Manager Author on Google Cloud Platform Google APIs and Translate API Using real-world examples, the authors first walk you through the basics of cloud computing, cloud terminologies and public cloud services. Then they dive right into Google Cloud Platform and how you can use it to tackle your challenges, build new products, analyze big data, and much more. Whether you're an independent developer, startup, or Fortune 500 company, you have never had easier access to world-class production, product development, and infrastructure tools. Google Cloud Platform is your ticket to leveraging your skills and knowledge into making reliable, scalable, and efficient products—just like how Google builds its own products.

This handbook covers a wide range of topics related to the collection, processing, analysis, and use of geospatial data in their various forms. This handbook provides an overview of how spatial computing technologies for big data can be organized and implemented to solve real-world problems. Diverse subdomains ranging from indoor mapping and navigation over trajectory computing to earth observation from space, are also present in this handbook. It combines fundamental contributions focusing on spatio-textual analysis, uncertain databases, and spatial statistics with application examples such as road network detection or colocation detection using GPUs. In summary, this handbook gives an essential introduction and overview of the rich field of spatial information science and big geospatial data. It introduces three different perspectives, which together define the field of big geospatial data: a societal, governmental, and governance perspective. It discusses questions of how the acquisition, distribution

and exploitation of big geospatial data must be organized both on the scale of companies and countries. A second perspective is a theory-oriented set of contributions on arbitrary spatial data with contributions introducing into the exciting field of spatial statistics or into uncertain databases. A third perspective is taking a very practical perspective to big geospatial data, ranging from chapters that describe how big geospatial data infrastructures can be implemented and how specific applications can be implemented on top of big geospatial data. This would include for example, research in historic map data, road network extraction, damage estimation from remote sensing imagery, or the analysis of spatio-textual collections and social media. This multi-disciplinary approach makes the book unique. This handbook can be used as a reference for undergraduate students, graduate students and researchers focused on big geospatial data. Professionals can use this book, as well as practitioners facing big collections of geospatial data.

Business activity monitoring, or BAM, provides real-time business intelligence by capturing data as it flows through a business system. By using BAM, you can monitor a business process in real time and generate alerts when the process needs human intervention. *Pro Business Activity Monitoring in BizTalk 2009* focuses on Microsoft's BAM tools, which provide a flexible infrastructure that captures data from Windows Communication Foundation, Windows Workflow Foundation, .NET applications, and BizTalk Server. This book shows why BAM is an important component of any business intelligence strategy because it bridges the gap between business intelligence and business process management. Part One of the book covers the basics of BAM. It teaches you how to install BAM, and then it leads you through a complete but simple project to capture data from a WCF service. Part Two of the book describes the BAM tool set. It shows you how to define the data that is captured by BAM, wire up a BizTalk application for BAM, and view BAM data using the BAM Portal. Part Three of the book focuses on monitoring applications written in WCF or WF. It also shows you how to use the BAM API. Part Four of the book covers advanced topics in BAM. These topics include integrating BAM with Microsoft's BI stack, creating relationships between BAM activities, managing BAM servers, and monitoring BizTalk applications using BAM. This book contains numerous exercises to help you learn BAM. Source code for the samples is available from Apress.com. For many CIOs, business intelligence is at the top of their priority lists. Learn why with *Pro Business Activity Monitoring in BizTalk 2009*.

API is technology and digital product used for artificial intelligence, platform economy, and internet. It has the capability to change business models dramatically. APIs (application programming interfaces) are becoming a major competitive factor for companies. This book takes on the fundamental questions of API Economy and approaches the subject pragmatically and clearly without technical jargon. The book clarifies the birth and shape of the API Economy with numerous practical examples. This is the first API Economy book based on scientific references. Originally this popular book was written in Finnish. It is a great start for students and advanced professionals alike. After reading this book, you will understand what it is all about and how to move forward and grow your business with APIs. The authors are leading Finnish API-experts with an abundance of experience from API and platform economy as authors, researchers, and lecturers and consultants.

This book offers a comprehensive introduction to relational (SQL) and non-relational (NoSQL) databases. The authors thoroughly review the current state of database tools and techniques, and examine coming innovations. The book opens with a broad look at data management,

including an overview of information systems and databases, and an explanation of contemporary database types: SQL and NoSQL databases, and their respective management systems The nature and uses of Big Data A high-level view of the organization of data management Data Modeling and Consistency Chapter-length treatment is afforded Data Modeling in both relational and graph databases, including enterprise-wide data architecture, and formulas for database design. Coverage of languages extends from an overview of operators, to SQL and and QBE (Query by Example), to integrity constraints and more. A full chapter probes the challenges of Ensuring Data Consistency, covering: Multi-User Operation Troubleshooting Consistency in Massive Distributed Data Comparison of the ACID and BASE consistency models, and more System Architecture also gets from its own chapter, which explores Processing of Homogeneous and Heterogeneous Data; Storage and Access Structures; Multi-dimensional Data Structures and Parallel Processing with MapReduce, among other topics. Post-Relational and NoSQL Databases The chapter on post-relational databases discusses the limits of SQL – and what lies beyond, including Multi-Dimensional Databases, Knowledge Bases and and Fuzzy Databases. A final chapter covers NoSQL Databases, along with Development of Non-Relational Technologies, Key-Value, Column-Family and Document Stores XML Databases and Graphical Databases, and more The book includes more than 100 tables, examples and illustrations, and each chapter offers a list of resources for further reading. SQL & NoSQL Databases conveys the strengths and weaknesses of relational and non-relational approaches, and shows how to undertake development for big data applications. The book benefits readers including students and practitioners working across the broad field of applied information technology. This textbook has been recommended and developed for university courses in Germany, Austria and Switzerland.

Explore and learn introductory topics about programming mechanisms for memory management available for Microsoft Windows. This book uses C++ pointers and specialized APIs such as the smart pointers of the C++ Standard Library and Microsoft UCRT functions. You'll also see how to work with lvalue and rvalue references. Introducing Mechanisms and APIs for Memory Management begins with topics about hardware features on the Intel x86 and Intel 64 (x64/amd64) hardware architectures and memory management. After reading this book you will be able to begin work with Windows memory management APIs. What You Will Learn Understand concepts and hardware features for Intel x86 and Intel 64 (x64/amd64) and memory management Discover C++ programming language techniques and smart pointers Work with Microsoft UCRT management APIs for memory management Who This Book Is For Software and cloud developers working on Microsoft Windows.

This book provides practical guidance for adopting a high velocity, continuous delivery process to create reliable, scalable, Software-as-a-Service (SaaS) solutions that are designed and built using a microservice architecture, deployed to the Azure cloud, and managed through automation. Microservices, IoT, and Azure offers software developers, architects, and operations engineers' step-by-step directions for building SaaS applications—applications that are available 24x7, work on any device, scale elastically, and are resilient to change--through code, script, exercises, and a working reference implementation. The book provides a working definition of microservices and contrasts this approach with traditional monolithic Layered Architecture. A fictitious, homebiomedical startup is used to demonstrate microservice architecture and automation capabilities for cross-cutting and business services as well as connected device scenarios for Internet of Things (IoT). Several Azure PaaS services are detailed including Storage, SQL Database, DocumentDb, Redis Cache, Cloud Services, Web API's, API Management, IoT Hub, IoT Suite, Event Hub, and Stream Analytics. Finally the book looks to the future and examines Service Fabric to see how microservices are becoming the de facto approach to building reliable software in the cloud. In this book, you'll learn: What

microservices are and why are they're a compelling architecture pattern for SaaS applications How to design, develop, and deploy microservices using Visual Studio, PowerShell, and Azure Microservice patterns for cross-cutting concerns and business capabilities Microservice patterns for Internet of Things and big data analytics solutions using IoT Hub, Event Hub, and Stream Analytics Techniques for automating microservice provisioning, building, and deployment What Service Fabric is and how it's the future direction for microservices on Microsoft Azure

This book explores various aspects of data engineering and information processing. In this second volume, the authors assess the challenges and opportunities involved in doing business with information. Their contributions on business information processing and management reflect diverse viewpoints – not only technological, but also business and social. As the global marketplace grows more and more complex due to the increasing availability of data, the information business is steadily gaining popularity and has a huge impact on modern society. Thus, there is a growing need for consensus on how business information can be created, accessed, used and managed. Discover the RESTful technologies, including REST, JSON, XML, JAX-RS web services, SOAP and more, for building today's microservices, big data applications, and web service applications. This book is based on a course the Oracle-based author is teaching for UC Santa Cruz Silicon Valley which covers architecture, design best practices and coding labs. Pro RESTful APIs: Design gives you all the fundamentals from the top down: from the top (architecture) through the middle (design) to the bottom (coding). This book is a must have for any microservices or web services developer building applications and services. What You'll Learn Discover the key RESTful APIs, including REST, JSON, XML, JAX, SOAP and more Use these for web services and data exchange, especially in today's big data context Harness XML, JSON, REST, and JAX-RS in examples and case studies Apply best practices to your solutions' architecture Who This Book Is For Experienced web programmers and developers.

In this book, cofounder and lead developer James Gardner brings you a comprehensive introduction to Pylons, the web framework that uses the best of Ruby, Python, and Perl and the emerging WSGI standard to provide structure and flexibility. You'll learn how to create your own Pylons-driven web site and attain the mastery of advanced Pylons features. You'll also learn how to stretch Pylons to its fullest ability, as well as share Gardner's unique insight and extensive experience in developing and deploying Pylons for a wide variety of situations. Beginning and experienced programmers will use this comprehensive guide to persistent memory programming. You will understand how persistent memory brings together several new software/hardware requirements, and offers great promise for better performance and faster application startup times—a huge leap forward in byte-addressable capacity compared with current DRAM offerings. This revolutionary new technology gives applications significant performance and capacity improvements over existing technologies. It requires a new way of thinking and developing, which makes this highly disruptive to the IT/computing industry. The full spectrum of industry sectors that will benefit from this technology include, but are not limited to, in-memory and traditional databases, AI, analytics, HPC, virtualization, and big data. Programming Persistent Memory describes the technology and why it is exciting the industry. It covers the operating system and hardware requirements as well as how to create development environments using emulated or real persistent memory hardware. The book explains fundamental concepts; provides an introduction to persistent memory programming APIs for C, C++, JavaScript, and other languages; discusses RMDA with persistent memory; reviews security features; and presents many examples. Source code and examples that you can run on your own systems are included. What You'll Learn Understand what persistent memory is, what it does, and the value it brings to the industry Become familiar with the operating system and hardware requirements to use persistent memory Know the fundamentals of

persistent memory programming: why it is different from current programming methods, and what developers need to keep in mind when programming for persistence Look at persistent memory application development by example using the Persistent Memory Development Kit (PMDK) Design and optimize data structures for persistent memory Study how real-world applications are modified to leverage persistent memory Utilize the tools available for persistent memory programming, application performance profiling, and debugging Who This Book Is For C, C++, Java, and Python developers, but will also be useful to software, cloud, and hardware architects across a broad spectrum of sectors, including cloud service providers, independent software vendors, high performance compute, artificial intelligence, data analytics, big data, etc.

Learn the business and technical importance of API design and architecture using the available cloud services from Azure and AWS. This book starts off with an introduction to APIs and the concept of API Economy from a business and organizational perspective. You'll decide on a sustainable API strategy and API architecture based on different case scenarios. You'll then look at actual examples on API development guidelines, providing a practical view and approach towards the API development and aligning teams in API development. This book walks you through the API gateway services available in Azure and AWS and reviews different approaches to API Security. This will prepare you for understanding the trade-off between security and the frictionless API experience. What You'll Learn Implement API Gateways to streamline API Development Examine Security Mapping with API gateways from Azure and AWS Apply API implementation using Serverless architecture Review evolving APIs for monitoring and changing business requirements Use code samples in API security implementations Who This Book Is For Developers and architects with .NET and web development experience who want to learn about API design.

A lot of work is required to release an API, but the effort doesn't always pay off. Overplanning before an API matures is a wasted investment, while underplanning can lead to disaster. This practical guide provides maturity models for individual APIs and multi-API landscapes to help you invest the right human and company resources for the right maturity level at the right time. How do you balance the desire for agility and speed with the need for robust and scalable operations? Four experts from the API Academy show software architects, program directors, and product owners how to maximize the value of their APIs by managing them as products through a continuous life cycle. Learn which API decisions you need to govern and how and where to do so Design, deploy, and manage APIs using an API-as-a-product (AaaP) approach Examine ten pillars that form the foundation of API product work Learn how the continuous improvement model governs changes throughout an API's lifetime Explore the five stages of a complete API product life cycle Delve into team roles needed to design, build, and maintain your APIs Learn how to manage your API landscape—the set of APIs published by your organization

E-book vendors continue to experiment: adjustments to business models, consolidation of content, and mergers with competitors mean constant change.

This book constitutes the thoroughly refereed post-conference proceedings of the Third COST Action IC1302 International KEYSTONE Conference on Semantic Keyword-Based Search on Structured Data Sources, IKC 2017, held in Gdańsk, Poland, in

September 2017. The 13 revised full papers and 5 short papers included in the first part of the book were carefully reviewed and selected from numerous submissions. The second part contains reports that summarize the major activities and achievements that have taken place in the context of the action: the short term scientific missions, the outcome of the summer schools, and the results achieved within the following four work packages: representation of structured data sources; keyword search; user interaction and keyword query interpretation; and research integration, showcases, benchmarks and evaluations. Also included is a short report generated by the chairs of the action. The papers cover a broad range of topics in the area of keyword search combining expertise from many different related fields such as information retrieval, natural language processing, ontology management, indexing, semantic web and linked data.

The computing technology on which we are now so dependent has risen to its position of ascendancy so rapidly that few of us have had the opportunity to take a step back and wonder where we are headed. This book urges us to do so. Taking a big-picture perspective on digital technology, *Living with Computers* leads the reader on a whistle-stop tour of the history of information and information technology. This journey culminates in a deep exploration into the meaning and role of computers in our lives, and what this experience might possibly mean for the future of human society – and the very existence of humanity itself. In the face of the transformative power of computing, this book provokes us to ask big questions. If computers become integrated into our bodies, merging with the information processing of our very DNA, will computing help to shape the evolution of biological life? If artificial intelligence advances beyond the abilities of the human brain, will this overturn our anthropocentrism and lead to a new view of reality? Will we control the computers of the future, or will they control us? These questions can be discomfiting, yet they cannot be ignored. This book argues that it is time to reshape our definition of our species in the context of our interaction with computing. For although such science-fiction scenarios are not likely to happen any time soon – and may, in fact, never happen – it is nevertheless vital to consider these issues now if we wish to have any influence over whatever is to come. So, humans, let's confront our possible destiny! James W. Cortada is a Senior Research Fellow at the Charles Babbage Institute at the University of Minnesota. He holds a Ph.D. in modern history and worked at IBM in various positions for 38 years, including in IBM's management research institute, The IBM Institute for Business Value (IBV). He is the author of over a dozen books on management, and nearly two dozen books on the history of information technology. These include the Springer title *From Urban Legends to Political Fact-Checking: Online Scrutiny in America, 1990-2015* (with William Aspray).

Learn how to accelerate C++ programs using data parallelism. This open access book enables C++ programmers to be at the forefront of this exciting and important new development that is helping to push computing to new levels. It is full of practical advice, detailed explanations, and code examples to illustrate key topics. Data parallelism in C++ enables access to parallel resources in a modern heterogeneous system, freeing you from being locked into any particular computing device. Now a single C++ application can use any combination of devices—including GPUs, CPUs, FPGAs and AI ASICs—that are suitable to the problems at hand. This book begins by introducing data parallelism and foundational topics for effective use of the SYCL standard

from the Khronos Group and Data Parallel C++ (DPC++), the open source compiler used in this book. Later chapters cover advanced topics including error handling, hardware-specific programming, communication and synchronization, and memory model considerations. Data Parallel C++ provides you with everything needed to use SYCL for programming heterogeneous systems. What You'll Learn Accelerate C++ programs using data-parallel programming Target multiple device types (e.g. CPU, GPU, FPGA) Use SYCL and SYCL compilers Connect with computing's heterogeneous future via Intel's oneAPI initiative Who This Book Is For Those new data-parallel programming and computer programmers interested in data-parallel programming using C++.

This book constitutes the proceedings of the 19th International Conference on Software and Systems Reuse, ICSR 2020, held in Hammamet, Tunisia in December 2020. Due to COVID-19 pandemic the Conference was held virtually. The 16 full papers and 2 short papers included in this book were carefully reviewed and selected from 60 submissions. The papers were organized in topical sections named: modelling, reuse in practice, reengineering, recommendation, and empirical analysis.

Implement application programming interface (API) usability, security, availability, reliability, and scalability to extend your company's market and potentially generate revenue. Businesses know they need to extend their markets into the digital world, and expose internal data to the Internet. This book shows how stakeholders within an organization can make it a successful journey. Stakeholder needs are not identical and departments experience difficulties discussing requirements with each other due to their different fundamental understanding of the process. The goal of this book is to introduce a common language for all business groups—developers, security experts, architects, product managers—around APIs and provide an overview of all aspects that need to be considered when exposing internal data. Most of the content in this book is based on feedback from real-world enterprise customer questions, challenges, and business scenarios. Practical guidance is provided on the business value of APIs, the general requirements to know, and how to undertake an audience-based implementation. You will learn how to protect access to data, as well as API error handling, documentation, management, integration, and more. What You'll Learn Know the types of APIs and their business and technical requirements The main benefits of APIs, including business value, loose coupling, and frequent updates Protect access to APIs through role-based access, attribute-based access, and rate limiting Distinguish between OAuth and OpenID Connect, and know how they both work Manage API error handling, including what should and should not be handled Understand the distinction between runtime, dynamic data, and static data Leverage external APIs as part of your own APIs Who This Book Is For API developers, API security experts, software architects, product owners, and business owners API Management Developing and Managing APIs for Your Organization Apress

This book provides a detailed review of how oncology drug development has changed over the past decade, and serves as a comprehensive guide for the practicalities in setting up phase I trials. The book covers strategies to accelerate the development of novel antitumor compounds from the laboratory to clinical trials and beyond through the use of innovative mechanism-of-action pharmacodynamic biomarkers and pharmacokinetic studies. The reader will learn about all aspects

of modern phase I trial designs, including the incorporation of precision medicine strategies, and approaches for rational patient allocation to novel anticancer therapies. Circulating biomarkers to assess mechanisms of response and resistance are changing the way we are assessing patient selection and are also covered in this book. The development of the different classes of antitumor agents are discussed, including chemotherapy, molecularly targeted agents, immunotherapies and also radiotherapy. The authors also discuss the lessons that the oncology field has learnt from the development of hematology-oncology drugs and how such strategies can be carried over into therapies for solid tumors. There is a dedicated chapter that covers the specialized statistical approaches necessary for phase I trial designs, including novel Bayesian strategies for dose escalation. This volume is designed to help clinicians better understand phase I clinical trials, but would also be of use to translational researchers (MDs and PhDs), and drug developers from academia and industry interested in cancer drug development. It could also be of use to phase I trial study coordinators, oncology nurses and advanced practice providers. Other health professionals interested in the treatment of cancer will also find this book of great value.

"This book offers suggestions, solutions, and recommendations for new and emerging research in Semantic Web technology, focusing broadly on methods and techniques for making the Web more useful and meaningful"--Provided by publisher.

Understand the key challenges and solutions around building microservices in the enterprise application environment. This book provides a comprehensive understanding of microservices architectural principles and how to use microservices in real-world scenarios. Architectural challenges using microservices with service integration and API management are presented and you learn how to eliminate the use of centralized integration products such as the enterprise service bus (ESB) through the use of composite/integration microservices. Concepts in the book are supported with use cases, and emphasis is put on the reality that most of you are implementing in a "brownfield" environment in which you must implement microservices alongside legacy applications with minimal disruption to your business. Microservices for the Enterprise covers state-of-the-art techniques around microservices messaging, service development and description, service discovery, governance, and data management technologies and guides you through the microservices design process. Also included is the importance of organizing services as core versus atomic, composite versus integration, and API versus edge, and how such organization helps to eliminate the use of a central ESB and expose services through an API gateway. What You'll Learn Design and develop microservices architectures with confidence Put into practice the most modern techniques around messaging technologies Apply the Service Mesh pattern to overcome inter-service communication challenges Apply battle-tested microservices security patterns to

address real-world scenarios Handle API management, decentralized data management, and observability Who This Book Is For Developers and DevOps engineers responsible for implementing applications around a microservices architecture, and architects and analysts who are designing such systems

This book describes the key concepts, principles and implementation options for creating high-assurance cloud computing solutions. The guide starts with a broad technical overview and basic introduction to cloud computing, looking at the overall architecture of the cloud, client systems, the modern Internet and cloud computing data centers. It then delves into the core challenges of showing how reliability and fault-tolerance can be abstracted, how the resulting questions can be solved, and how the solutions can be leveraged to create a wide range of practical cloud applications. The author's style is practical, and the guide should be readily understandable without any special background. Concrete examples are often drawn from real-world settings to illustrate key insights. Appendices show how the most important reliability models can be formalized, describe the API of the Isis2 platform, and offer more than 80 problems at varying levels of difficulty.

This book constitutes selected papers from the 16th European, Mediterranean, and Middle Eastern Conference, EMCIS 2019, held in Dubai, UAE, in October 2019. EMCIS is dedicated to the definition and establishment of Information Systems as a discipline of high impact for the methodical community and IS professionals, focusing on approaches that facilitate the identification of innovative research of significant relevance to the IS discipline. The 48 full papers presented in this volume were carefully reviewed and selected from a total of 138 submissions. They were organized in topical sections named: Big Data and Analytics; Blockchain Technology and Applications; Cloud Computing; Digital Services and Social Media; e-Government; Enterprise Information Systems; Health-Care Information Systems; Information Systems Security and Information Privacy Protection; Innovative Research Projects; IT Governance; and Management and Organizational Issues in Information Systems.

This textbook begins with an introduction to the US healthcare delivery system, its many systemic challenges and the prior efforts to develop and deploy informatics tools to help overcome those problems. It goes on to discuss health informatics from an historical perspective, its current state and its likely future state now that electronic health record systems are widely deployed, the HL7 Fast Healthcare Interoperability standard is being rapidly accepted as the means to access the data stored in those systems and analytics is increasing being used to gain new knowledge from that aggregated clinical data. It then turns to some of the important and evolving areas of informatics including population and public health, mHealth and big data and analytics. Use cases and case studies are used in all of these discussions to help readers connect the technologies to real world challenges. Effective use of informatics systems and tools by

providers and their patients is key to improving the quality, safety and cost of healthcare. With health records now digital, no effective means has existed for sharing them with patients, among the multiple providers who may care for them and for important secondary uses such as public/population health and research. This problem is a topic of congressional discussion and is addressed by the 21st Century Cures Act of 2016 that mandates that electronic health record (EHR) systems offer a patient-facing API. HL7's Fast Healthcare Interoperability Resources (FHIR) is that API and this is the first comprehensive treatment of the technology and the many ways it is already being used. FHIR is based on web technologies and is thus a far more facile, easy to implement approach that is rapidly gaining acceptance. It is also the basis for a 'universal health app platform' that literally has the potential to foster innovation around the data in patient records similar to the app ecosystems smartphones created around the data they store. FHIR app stores have already been opened by Epic and Cerner, the two largest enterprise EHR vendors. Provider facing apps are already being explored to improve EHR usability and support personalized medicine. Medicare and the Veteran's Administration have announced FHIR app platforms for their patients. Apple's new IOS 11.3 features the ability for consumers to aggregate their health records on their iPhone using FHIR. Health insurance companies are exploring applications of FHIR to improve service and communication with their providers and patients. SureScripts, the national e-Prescribing network, is using FHIR to help doctors know if their patients are complying with prescriptions. This textbook is for introductory health informatics courses for computer science and health sciences students (e.g. doctors, nurses, PhDs), the current health informatics community, IT professionals interested in learning about the field and practicing healthcare providers. Though this textbook covers an important new technology, it is accessible to non-technical readers including healthcare providers, their patients or anyone interested in the use of healthcare data for improved care, public/population health or research.

[Copyright: 644b32f874031b22ad3232d6e5af3197](#)