

Distributed And Cloud Computing 1st Edition Elsevier

This book gathers research contributions on recent advances in intelligent and distributed computing. A major focus is placed on new techniques and applications for several highly demanded research directions: Internet of Things, Cloud Computing and Big Data, Data Mining and Machine Learning, Multi-agent and Service-Based Distributed Systems, Distributed Algorithms and Optimization, Modeling Operational Processes, Social Network Analysis and Inappropriate Content Counteraction, Cyber-Physical Security and Safety, Intelligent Distributed Decision Support Systems, Intelligent Human-Machine Interfaces, Visual Analytics and others. The book represents the peer-reviewed proceedings of the 13th International Symposium on Intelligent Distributed Computing (IDC 2019), which was held in St. Petersburg, Russia, from October 7 to 9, 2019.

Computer architecture deals with the physical configuration, logical structure, formats, protocols, and operational sequences for processing data, controlling the configuration, and controlling the operations over a computer. It also encompasses word lengths, instruction codes, and the interrelationships among the main parts of a computer or group of computers. This two-volume set offers a comprehensive coverage of the field of computer organization and architecture.

Many applications follow the distributed computing paradigm, in which parts of the application are executed on different network-interconnected computers. The extension of these applications in terms of number of users or size has led to an unprecedented increase in the scale of the infrastructure that supports them. Large-Scale Distributed Computing and Applications: Models and Trends offers a coherent and realistic image of today's research results in large scale distributed systems, explains state-of-the-art technological solutions for the main issues regarding large scale distributed systems, and presents the benefits of using large scale distributed systems and the development process of scientific and commercial distributed applications.

Mastering Cloud Computing is designed for undergraduate students learning to develop cloud computing applications. Tomorrow's applications won't live on a single computer but will be deployed from and reside on a virtual server, accessible anywhere, any time. Tomorrow's application developers need to understand the requirements of building apps for these virtual systems, including concurrent programming, high-performance computing, and data-intensive systems. The book introduces the principles of distributed and parallel computing underlying cloud architectures and specifically focuses on virtualization, thread programming, task programming, and map-reduce programming. There are examples demonstrating all of these and more, with exercises and labs throughout. Explains how to make design choices and tradeoffs to consider when building applications to run in a virtual cloud environment Real-world case studies include scientific, business, and energy-efficiency considerations

Distributed Computing Through Combinatorial Topology describes techniques for analyzing distributed algorithms based on award winning combinatorial topology research. The authors present a solid theoretical foundation relevant to many real systems reliant on parallelism with unpredictable delays, such as multicore microprocessors, wireless networks, distributed systems, and Internet protocols. Today, a new student or researcher must assemble a collection of scattered conference publications, which are typically terse and commonly use different notations and terminologies. This book provides a self-contained explanation of the mathematics to readers with computer science backgrounds, as well as explaining computer science concepts to readers with backgrounds in applied mathematics. The first section presents mathematical notions and models, including message passing and shared-memory systems, failures, and timing models. The next section presents core concepts in two chapters each: first, proving a simple result that lends itself to examples and pictures that will build up readers' intuition; then generalizing the concept to prove a more sophisticated result. The overall result weaves together and develops the basic concepts of the field, presenting them in a gradual and intuitively appealing way. The book's final section discusses advanced topics typically found in a graduate-level course for those who wish to explore further. Named a 2013 Notable Computer Book for Computing Methodologies by Computing Reviews Gathers knowledge otherwise spread across research and conference papers using consistent notations and a standard approach to facilitate understanding Presents unique insights applicable to multiple computing fields, including multicore microprocessors, wireless networks, distributed systems, and Internet protocols Synthesizes and distills material into a simple, unified presentation with examples, illustrations, and exercises

Global networks, which are the primary pillars of the modern manufacturing industry and supply chains, can only cope with the new challenges, requirements and demands when supported by new computing and Internet-based technologies. Cloud Manufacturing: Distributed Computing Technologies for Global and Sustainable Manufacturing introduces a new paradigm for scalable service-oriented sustainable and globally distributed manufacturing systems. The eleven chapters in this book provide an updated overview of the latest technological development and applications in relevant research areas. Following an introduction to the essential features of Cloud Computing, chapters cover a range of methods and applications such as the factors that actually affect adoption of the Cloud Computing technology in manufacturing companies and new geometrical simplification method to stream 3-Dimensional design and manufacturing data via the Internet. This is further supported case studies and real life data for Waste Electrical and Electronic Equipment (WEEE) remanufacturing. This compilation of up to date research and literature can be used as a textbook or reference for mechanical, manufacturing, and computer engineering graduate students and researchers for efficient utilization, deployment and development of distributed and Cloud manufacturing systems, services and applications.

The primary purpose of this book is to capture the state-of-the-art in Cloud Computing technologies and applications. The book will also aim to identify potential research directions and technologies that will facilitate creation a global market-place of cloud computing services supporting scientific, industrial, business, and consumer applications. We expect the book to serve as a reference for larger audience such as systems architects, practitioners, developers, new researchers and graduate level students. This area of research is relatively recent, and as such has no existing reference book that addresses it. This book will be a timely contribution to a field that is gaining considerable research interest, momentum, and is expected to be of increasing interest to commercial developers. The book is targeted for professional computer science developers and graduate students especially at Masters level. As Cloud Computing is recognized as one of the top five emerging technologies that will have a major impact on the quality of science and society over the next 20 years, its knowledge will help position our readers at the forefront of the field.

Comprehensive and timely, *Cloud Computing: Concepts and Technologies* offers a thorough and detailed description of cloud computing concepts, architectures, and technologies, along with guidance on the best ways to understand and implement them. It covers the multi-core architectures, distributed and parallel computing models, virtualization, cloud developments, workload and Service-Level-Agreements (SLA) in cloud, workload management. Further, resource management issues in cloud with regard to resource provisioning, resource allocation, resource mapping and resource adaptation, ethical, non-ethical and security issues in cloud are followed by discussion of open challenges and future directions. This book gives students a comprehensive overview of the latest technologies and guidance on cloud computing, and is ideal for those studying the subject in specific modules or advanced courses. It is designed in twelve chapters followed by laboratory setups and experiments. Each chapter has multiple choice questions with answers, as well as review questions and critical thinking questions. The chapters are practically-focused, meaning that the information will also be relevant and useful for professionals wanting an overview of the topic.

“There’s an incredible amount of depth and thinking in the practices described here, and it’s impressive to see it all in one place.” —Win Treese, coauthor of *Designing Systems for Internet Commerce* The *Practice of Cloud System Administration, Volume 2*, focuses on “distributed” or “cloud” computing and brings a DevOps/SRE sensibility to the practice of system administration. Unsatisfied with books that cover either design or operations in isolation, the authors created this authoritative reference centered on a comprehensive approach. Case studies and examples from Google, Etsy, Twitter, Facebook, Netflix, Amazon, and other industry giants are explained in practical ways that are useful to all enterprises. The new companion to the best-selling first volume, *The Practice of System and Network Administration, Second Edition*, this guide offers expert coverage of the following and many other crucial topics: Designing and building modern web and distributed systems Fundamentals of large system design Understand the new software engineering implications of cloud administration Make systems that are resilient to failure and grow and scale dynamically Implement DevOps principles and cultural changes IaaS/PaaS/SaaS and virtual platform selection Operating and running systems using the latest DevOps/SRE strategies Upgrade production systems with zero down-time What and how to automate; how to decide what not to automate On-call best practices that improve uptime Why distributed systems require fundamentally different system administration techniques Identify and resolve resiliency problems before they surprise you Assessing and evaluating your team’s operational effectiveness Manage the scientific process of continuous improvement A forty-page, pain-free assessment system you can start using today

Master Powerful Off-the-Shelf Business Solutions for AI and Machine Learning Pragmatic AI will help you solve real-world problems with contemporary machine learning, artificial intelligence, and cloud computing tools. Noah Gift demystifies all the concepts and tools you need to get results—even if you don’t have a strong background in math or data science. Gift illuminates powerful off-the-shelf cloud offerings from Amazon, Google, and Microsoft, and demonstrates proven techniques using the Python data science ecosystem. His workflows and examples help you streamline and simplify every step, from deployment to production, and build exceptionally scalable solutions. As you learn how machine language (ML) solutions work, you’ll gain a more intuitive understanding of what you can achieve with them and how to maximize their value. Building on these fundamentals, you’ll walk step-by-step through building cloud-based AI/ML applications to address realistic issues in sports marketing, project management, product pricing, real estate, and beyond. Whether you’re a business professional, decision-maker, student, or programmer, Gift’s expert guidance and wide-ranging case studies will prepare you to solve data science problems in virtually any environment. Get and configure all the tools you’ll need Quickly review all the Python you need to start building machine learning applications Master the AI and ML toolchain and project lifecycle Work with Python data science tools such as IPython, Pandas, Numpy, Jupyter Notebook, and Sklearn Incorporate a pragmatic feedback loop that continually improves the efficiency of your workflows and systems Develop cloud AI solutions with Google Cloud Platform, including TPU, Colaboratory, and Datalab services Define Amazon Web Services cloud AI workflows, including spot instances, code pipelines, boto, and more Work with Microsoft Azure AI APIs Walk through building six real-world AI applications, from start to finish Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

This book contains a selection of refereed and revised papers of the Intelligent Distributed Computing Track originally presented at the third International Symposium on Intelligent Informatics (ISI-2014), September 24-27, 2014, Delhi, India. The papers selected for this Track cover several Distributed Computing and related topics including Peer-to-Peer Networks, Cloud Computing, Mobile Clouds, Wireless Sensor Networks, and their applications.

Explores cloud computing, breaking down the concepts, models, mechanisms, and architectures of this technology while allowing for the financial assessment of resources and how they compare to traditional storage systems.

Mobile Cloud Computing: Foundations and Service Models combines cloud computing, mobile computing and wireless networking to bring new computational resources for mobile users, network operators and cloud computing providers. The book provides the latest research and development insights on mobile cloud computing, beginning with an exploration of the foundations of cloud computing, existing cloud infrastructures classifications, virtualization techniques and service models. It then examines the approaches to building cloud services using a bottom-up approach, describing data center design, cloud networking and software orchestration solutions, showing how these solutions support mobile devices and services. The book describes mobile cloud clouding concepts with a particular focus on a user-centric approach, presenting a distributed mobile cloud service model called POEM to manage mobile cloud resource and compose mobile cloud applications. It concludes with a close examination of the security and privacy issues of mobile clouds. Shows how to construct new mobile cloud based applications Contains detailed approaches to address security challenges in mobile cloud computing Includes a case

study using vehicular cloud

Cloud computing offers subscription-based on-demand services, and it has emerged as the backbone of the computing industry. It has enabled us to share resources among multiple users through virtualization, which creates a virtual instance of a computer system running in an abstracted hardware layer. Unlike early distributed computing models, it offers virtually limitless computing resources through its large scale cloud data centers. It has gained wide popularity over the past few years, with an ever-increasing infrastructure, a number of users, and the amount of hosted data. The large and complex workloads hosted on these data centers introduce many challenges, including resource utilization, power consumption, scalability, and operational cost. Therefore, an effective resource management scheme is essential to achieve operational efficiency with improved elasticity. Machine learning enabled solutions are the best fit to address these issues as they can analyze and learn from the data. Moreover, it brings automation to the solutions, which is an essential factor in dealing with large distributed systems in the cloud paradigm. Machine Learning for Cloud Management explores cloud resource management through predictive modelling and virtual machine placement. The predictive approaches are developed using regression-based time series analysis and neural network models. The neural network-based models are primarily trained using evolutionary algorithms, and efficient virtual machine placement schemes are developed using multi-objective genetic algorithms. Key Features: The first book to set out a range of machine learning methods for efficient resource management in a large distributed network of clouds. Predictive analytics is an integral part of efficient cloud resource management, and this book gives a future research direction to researchers in this domain. It is written by leading international researchers. The book is ideal for researchers who are working in the domain of cloud computing.

A tutorial leading the aspiring Go developer to full mastery of Golang's distributed features. Key Features This book provides enough concurrency theory to give you a contextual understanding of Go concurrency It gives weight to synchronous and asynchronous data streams in Golang web applications It makes Goroutines and Channels completely familiar and natural to Go developers Book Description Distributed Computing with Go gives developers with a good idea how basic Go development works the tools to fulfill the true potential of Golang development in a world of concurrent web and cloud applications. Nikhil starts out by setting up a professional Go development environment. Then you'll learn the basic concepts and practices of Golang concurrent and parallel development. You'll find out in the new few chapters how to balance resources and data with REST and standard web approaches while keeping concurrency in mind. Most Go applications these days will run in a data center or on the cloud, which is a condition upon which the next chapter depends. There, you'll expand your skills considerably by writing a distributed document indexing system during the next two chapters. This system has to balance a large corpus of documents with considerable analytical demands. Another use case is the way in which a web application written in Go can be consciously redesigned to take distributed features into account. The chapter is rather interesting for Go developers who have to migrate existing Go applications to computationally and memory-intensive environments. The final chapter relates to the rather onerous task of testing parallel and distributed applications, something that is not usually taught in standard computer science curricula. What you will learn Gain proficiency with concurrency and parallelism in Go Learn how to test your application using Go's standard library Learn industry best practices with technologies such as REST, OpenAPI, Docker, and so on Design and build a distributed search engine Learn strategies on how to design a system for web scale Who this book is for This book is for developers who are familiar with the Golang syntax and have a good idea of how basic Go development works. It would be advantageous if you have been through a web application product cycle, although it's not necessary.

Distributed and Cloud Computing: From Parallel Processing to the Internet of Things offers complete coverage of modern distributed computing technology including clusters, the grid, service-oriented architecture, massively parallel processors, peer-to-peer networking, and cloud computing. It is the first modern, up-to-date distributed systems textbook; it explains how to create high-performance, scalable, reliable systems, exposing the design principles, architecture, and innovative applications of parallel, distributed, and cloud computing systems. Topics covered by this book include: facilitating management, debugging, migration, and disaster recovery through virtualization; clustered systems for research or ecommerce applications; designing systems as web services; and social networking systems using peer-to-peer computing. The principles of cloud computing are discussed using examples from open-source and commercial applications, along with case studies from the leading distributed computing vendors such as Amazon, Microsoft, and Google. Each chapter includes exercises and further reading, with lecture slides and more available online. This book will be ideal for students taking a distributed systems or distributed computing class, as well as for professional system designers and engineers looking for a reference to the latest distributed technologies including cloud, P2P and grid computing. Complete coverage of modern distributed computing technology including clusters, the grid, service-oriented architecture, massively parallel processors, peer-to-peer networking, and cloud computing Includes case studies from the leading distributed computing vendors: Amazon, Microsoft, Google, and more Explains how to use virtualization to facilitate management, debugging, migration, and disaster recovery Designed for undergraduate or graduate students taking a distributed systems course—each chapter includes exercises and further reading, with lecture slides and more available online

Cloud computing continues to emerge as a subject of substantial industrial and academic interest. Although the meaning and scope of “cloud computing” continues to be debated, the current notion of clouds blurs the distinctions between grid services, web services, and data centers, among other areas. Clouds also bring considerations of lowering the cost for relatively bursty applications to the fore. Cloud Computing: Principles, Systems and Applications is an essential reference/guide that provides thorough and timely examination of the services, interfaces and types of applications that can be executed on cloud-based systems. The book identifies and highlights state-of-the-art techniques and methods for designing cloud systems, presents

mechanisms and schemes for linking clouds to economic activities, and offers balanced coverage of all related technologies that collectively contribute towards the realization of cloud computing. With an emphasis on the conceptual and systemic links between cloud computing and other distributed computing approaches, this text also addresses the practical importance of efficiency, scalability, robustness and security as the four cornerstones of quality of service. Topics and features: explores the relationship of cloud computing to other distributed computing paradigms, namely peer-to-peer, grids, high performance computing and web services; presents the principles, techniques, protocols and algorithms that can be adapted from other distributed computing paradigms to the development of successful clouds; includes a Foreword by Professor Mark Baker of the University of Reading, UK; examines current cloud-practical applications and highlights early deployment experiences; elaborates the economic schemes needed for clouds to become viable business models. This book will serve as a comprehensive reference for researchers and students engaged in cloud computing. Professional system architects, technical managers, and IT consultants will also find this unique text a practical guide to the application and delivery of commercial cloud services. Prof. Nick Antonopoulos is Head of the School of Computing, University of Derby, UK. Dr. Lee Gillam is a Lecturer in the Department of Computing at the University of Surrey, UK.

Fueled by ubiquitous computing ambitions, the edge is at the center of confluence of many emergent technological trends such as hardware-rooted trust and code integrity, 5G, data privacy and sovereignty, blockchains and distributed ledgers, ubiquitous sensors and drones, autonomous systems and real-time stream processing. Hardware and software pattern maturity have reached a tipping point so that scenarios like smart homes, smart factories, smart buildings, smart cities, smart grids, smart cars, smart highways are in reach of becoming a reality. While there is a great desire to bring born-in-the-cloud patterns and technologies such as zero-downtime software and hardware updates/upgrades to the edge, developers and operators alike face a unique set of challenges due to environmental differences such as resource constraints, network availability and heterogeneity of the environment. The first part of the book discusses various edge computing patterns which the authors have observed, and the reasons why these observations have led them to believe that there is a need for a new architectural paradigm for the new problem domain. Edge computing is examined from the app designer and architect's perspectives. When they design for edge computing, they need a new design language that can help them to express how capabilities are discovered, delivered and consumed, and how to leverage these capabilities regardless of location and network connectivity. Capability-Oriented Architecture is designed to provide a framework for all of these. This book is for everyone who is interested in understanding what ubiquitous and edge computing means, why it is growing in importance and its opportunities to you as a technologist or decision maker. The book covers the broad spectrum of edge environments, their challenges and how you can address them as a developer or an operator. The book concludes with an introduction to a new architectural paradigm called capability-based architecture, which takes into consideration the capabilities provided by an edge environment. .

This book introduces the reader to the fundamentals of contemporary, emerging and future technologies and services in Internet computing. It covers essential concepts such as distributed systems architectures and web technologies, contemporary paradigms such as cloud computing and the Internet of things, and emerging technologies like distributed ledger technologies and fog computing. The book also highlights the interconnection and recombination of these Internet-based technologies, which together form a critical information infrastructure with major impacts on individuals, organizations, governments, economies, and society as a whole. Intended as a textbook for upper undergraduate and graduate classes, it features a wealth of examples, learning goals and summaries for every chapter, numerous recommendations for further reading, and questions for checking students' comprehension. A dedicated author website offers additional teaching material and more elaborate examples. Accordingly, the book enables students and young professionals in IT-related fields to familiarize themselves with the Internet's basic mechanisms, and with the most promising Internet-based technologies of our time.

This book provides a technical description of cloud computing technologies, covering cloud infrastructure and platform services. It then addresses the basics of operating a Cloud computing data center, the services offered from Cloud providers, the carrier role in connecting users to data centers, and the process of interconnecting Cloud data centers to form a flexible processing unit. It also describes how cloud computing has made an impact in various industries and provides emerging technologies that are critical within each industry. Lastly, this book will address security requirements and provide the best practices in securing data.

This volume contains the proceedings of CloudCom 2009, the First International Conference on Cloud Computing. The conference was held in Beijing, China, during December 1–4, 2009, and was the first in a series initiated by the Cloud Computing Association (www.cloudcom.org). The Cloud Computing Association was founded in 2009 by Chunming Rong, Martin Gilje Jaatun, and Frode Eika Sandnes. This first conference was organized by the Beijing Jitong University, Chinese Institute of Electronics, and Wuhan University, and co-organized by Huazhong University of Science and Technology, South China Normal University, and Sun Yat-sen University. Ever since the inception of the Internet, a "Cloud" has been used as a metaphor for a network-accessible infrastructure (e.g., data storage, computing hardware, or entire networks) which is hidden from users. To some, the concept of cloud computing may seem like a throwback to the days of big mainframe computers, but we believe that cloud computing makes data truly mobile, allowing a user to access services anywhere, anytime, with any Internet browser. In cloud computing, IT-related capabilities are provided as services, accessible without requiring control of, or even knowledge of, the underlying technology. Cloud computing provides dynamic scalability of services and computing power, and although many mature technologies are used as components in cloud computing, there are still many unresolved and open problems.

Distributed and Cloud Computing: From Parallel Processing to the Internet of Things offers complete coverage of modern distributed computing technology including clusters, the grid, service-oriented architecture, massively parallel processors, peer-to-peer networking, and cloud computing. It is the first modern, up-to-date distributed systems textbook; it explains how to create high-performance, scalable, reliable systems, exposing the design principles, architecture, and innovative applications of parallel, distributed, and cloud computing systems. Topics covered by this book include: facilitating management, debugging, migration, and disaster recovery through virtualization; clustered systems for research or ecommerce applications; designing systems as web services; and social networking systems using peer-to-peer computing. The principles of cloud computing are discussed using examples from open-source and commercial applications,

along with case studies from the leading distributed computing vendors such as Amazon, Microsoft, and Google. Each chapter includes exercises and further reading, with lecture slides and more available online. This book will be ideal for students taking a distributed systems or distributed computing class, as well as for professional system designers and engineers looking for a reference to the latest distributed technologies including cloud, P2P and grid computing. Complete coverage of modern distributed computing technology including clusters, the grid, service-oriented architecture, massively parallel processors, peer-to-peer networking, and cloud computing Includes case studies from the leading distributed computing vendors: Amazon, Microsoft, Google, and more Explains how to use virtualization to facilitate management, debugging, migration, and disaster recovery Designed for undergraduate or graduate students taking a distributed systems course--each chapter includes exercises and further reading, with lecture slides and more available online.

The Encyclopedia of Cloud Computing provides IT professionals, educators, researchers and students with a compendium of cloud computing knowledge. Authored by a spectrum of subject matter experts in industry and academia, this unique publication, in a single volume, covers a wide range of cloud computing topics, including technological trends and developments, research opportunities, best practices, standards, and cloud adoption. Providing multiple perspectives, it also addresses questions that stakeholders might have in the context of development, operation, management, and use of clouds. Furthermore, it examines cloud computing's impact now and in the future. The encyclopedia presents 56 chapters logically organized into 10 sections. Each chapter covers a major topic/area with cross-references to other chapters and contains tables, illustrations, side-bars as appropriate. Furthermore, each chapter presents its summary at the beginning and backend material, references and additional resources for further information.

This book provides the users with quick and easy data acquisition, processing, storage and product generation services. It describes the entire life cycle of remote sensing data and builds an entire high performance remote sensing data processing system framework. It also develops a series of remote sensing data management and processing standards. Features: Covers remote sensing cloud computing Covers remote sensing data integration across distributed data centers Covers cloud storage based remote sensing data share service Covers high performance remote sensing data processing Covers distributed remote sensing products analysis

Guide to Cloud Computing for Business and Technology Managers: From Distributed Computing to Cloudware Applications unravels the mystery of cloud computing and explains how it can transform the operating contexts of business enterprises. It provides a clear understanding of what cloud computing really means, what it can do, and when it is practical

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.

Cloud Computing: Theory and Practice provides students and IT professionals with an in-depth analysis of the cloud from the ground up. Beginning with a discussion of parallel computing and architectures and distributed systems, the book turns to contemporary cloud infrastructures, how they are being deployed at leading companies such as Amazon, Google and Apple, and how they can be applied in fields such as healthcare, banking and science. The volume also examines how to successfully deploy a cloud application across the enterprise using virtualization, resource management and the right amount of networking support, including content delivery networks and storage area networks. Developers will find a complete introduction to application development provided on a variety of platforms. Learn about recent trends in cloud computing in critical areas such as: resource management, security, energy consumption, ethics, and complex systems Get a detailed hands-on set of practical recipes that help simplify the deployment of a cloud based system for practical use of computing clouds along with an in-depth discussion of several projects Understand the evolution of cloud computing and why the cloud computing paradigm has a better chance to succeed than previous efforts in large-scale distributed computing

A comprehensive guide to Fog and Edge applications, architectures, and technologies Recent years have seen the explosive growth of the Internet of Things (IoT): the internet-connected network of devices that includes everything from personal electronics and home appliances to automobiles and industrial machinery. Responding to the ever-increasing bandwidth demands of the IoT, Fog and Edge computing concepts have developed to collect, analyze, and process data more efficiently than traditional cloud architecture. Fog and Edge Computing: Principles and Paradigms provides a comprehensive overview of the state-of-the-art applications and architectures driving this dynamic field of computing while highlighting potential research directions and emerging technologies. Exploring topics such as developing scalable architectures, moving from closed systems to open systems, and ethical issues rising from data sensing, this timely book addresses both the challenges and opportunities that Fog and Edge computing presents. Contributions from leading IoT experts discuss federating Edge resources, middleware design issues, data management and predictive analysis, smart transportation and surveillance applications, and more. A coordinated and integrated presentation of topics helps readers gain thorough knowledge of the foundations, applications, and issues that are central to Fog and Edge computing. This valuable resource: Provides insights on transitioning from current Cloud-centric and 4G/5G wireless environments to Fog Computing Examines methods to optimize virtualized, pooled, and shared resources Identifies potential technical challenges and offers suggestions for possible solutions Discusses major components of Fog and Edge computing architectures such as middleware, interaction protocols, and autonomic management Includes access to a

website portal for advanced online resources Fog and Edge Computing: Principles and Paradigms is an essential source of up-to-date information for systems architects, developers, researchers, and advanced undergraduate and graduate students in fields of computer science and engineering.

The Cloud is an advanced and fast-growing technology in the current era. The computing paradigm has changed drastically. It provided a new insight into the computing world with new characteristics including on-demand, virtualization, scalability and many more. Utility computing, virtualization and service-oriented architecture (SoA) are the key characteristics of Cloud computing. The Cloud provides distinct IT services over the web on a pay-as-you-go and on-demand basis. Cloud Computing Technologies for Smart Agriculture and Healthcare covers Cloud management and its framework. It also focuses how the Cloud computing framework can be integrated with applications based on agriculture and healthcare. Features: Contains a systematic overview of the state-of-the-art, basic theories, challenges, implementation, and case studies on Cloud technology Discusses of recent research results and future advancement in virtualization technology Focuses on core theories, architectures, and technologies necessary to develop and understand the computing models and its applications Includes a wide range of examples that uses Cloud technology for increasing farm profitability and sustainable production Presents the farming industry with Cloud technology that allows it to aggregate, analyze, and share data across farms and the world Includes Cloud-based electronic health records with privacy and security features Offers suitable IT solutions to the global issues in the domain of agriculture and health care for society This reference book is aimed at undergraduate and post-graduate programs. It will also help research scholars in their research work. This book also benefits like scientists, business innovators, entrepreneurs, professionals, and practitioners.

"This book focuses on network management and traffic engineering for Internet and distributed computing technologies, as well as present emerging technology trends and advanced platforms"--Provided by publisher.

AI and Cloud Computing, Volume 120 in the Advances in Computers series, highlights new advances in the field, with this updated volume presenting interesting chapters on topics including A Deep-forest based Approach for Detecting Fraudulent Online Transaction, Design of Cyber-Physical-Social Systems with Forensic-awareness Based on Deep Learning, Review on Privacy-preserving Data Comparison Protocols in Cloud Computing, Fingerprint Liveness Detection Using an Improved CNN with the Spatial Pyramid Pooling Structure, Protecting Personal Sensitive Data Security in the Cloud with Blockchain, and more. Provides the authority and expertise of leading contributors from an international board of authors Presents the latest release in the Advances in Computers series Includes the latest information on AI and Cloud Computing

The book focuses on mobile agents, which are computer programs that can autonomously migrate between network sites. This text introduces the concepts and principles of mobile agents, provides an overview of mobile agent technology, and focuses on applications in networking and distributed computing.

"This book compiles recent research trends and practical issues in the fields of distributed computing and Internet technologies, providing advancements on emerging technologies that aim to support the effective design and implementation of service-oriented networks, future Internet environments and building management frameworks"--

Concurrent and Distributed Computing in Java addresses fundamental concepts in concurrent computing with Java examples. The book consists of two parts. The first part deals with techniques for programming in shared-memory based systems. The book covers concepts in Java such as threads, synchronized methods, waits, and notify to expose students to basic concepts for multi-threaded programming. It also includes algorithms for mutual exclusion, consensus, atomic objects, and wait-free data structures. The second part of the book deals with programming in a message-passing system. This part covers resource allocation problems, logical clocks, global property detection, leader election, message ordering, agreement algorithms, checkpointing, and message logging. Primarily a textbook for upper-level undergraduates and graduate students, this thorough treatment will also be of interest to professional programmers. Designing distributed computing systems is a complex process requiring a solid understanding of the design problems and the theoretical and practical aspects of their solutions. This comprehensive textbook covers the fundamental principles and models underlying the theory, algorithms and systems aspects of distributed computing. Broad and detailed coverage of the theory is balanced with practical systems-related issues such as mutual exclusion, deadlock detection, authentication, and failure recovery. Algorithms are carefully selected, lucidly presented, and described without complex proofs. Simple explanations and illustrations are used to elucidate the algorithms. Important emerging topics such as peer-to-peer networks and network security are also considered. With vital algorithms, numerous illustrations, examples and homework problems, this textbook is suitable for advanced undergraduate and graduate students of electrical and computer engineering and computer science. Practitioners in data networking and sensor networks will also find this a valuable resource. Additional resources are available online at www.cambridge.org/9780521876346.

Cloud computing-accessing computing resources over the Internet-is rapidly changing the landscape of information technology. Its primary benefits compared to on-premise computing models are reduced costs and increased agility and scalability. Hence, cloud computing is receiving considerable interest among several stakeholders-businesses, the IT ind

Guide to Cloud Computing for Business and Technology Managers: From Distributed Computing to Cloudware Applications unravels the mystery of cloud computing and explains how it can transform the operating contexts of business enterprises. It provides a clear understanding of what cloud computing really means, what it can do, and when it is practical to use. Addressing the primary management and operation concerns of cloudware, including performance, measurement, monitoring, and security, this pragmatic book: Introduces the enterprise applications integration (EAI) solutions that were a first step toward enabling an integrated enterprise Details service-oriented architecture (SOA) and related technologies that paved the road for cloudware applications Covers delivery models like IaaS, PaaS, and SaaS, and deployment models like public, private, and hybrid clouds Describes Amazon, Google, and Microsoft cloudware solutions and services, as well as those of several other players Demonstrates how cloud computing can reduce costs, achieve business flexibility, and sharpen strategic focus Unlike customary discussions of cloud computing, Guide to Cloud Computing for Business and Technology Managers: From Distributed Computing to Cloudware Applications emphasizes the key differentiator—that cloud computing is able to treat enterprise-level services not merely as discrete stand-alone services, but as Internet-locatable, composable, and repackagable building blocks for generating dynamic real-world enterprise business processes.

“Computing Networks” explores the core of the new distributed computing infrastructures we are using today: the networking systems of clusters, grids and clouds. It helps network designers and distributed-application developers and users to better understand the technologies, specificities, constraints and benefits of these different infrastructures’ communication systems. Cloud Computing will give the possibility for millions of users to process data anytime, anywhere, while being eco-friendly. In order to deliver this emerging traffic in a timely, cost-efficient, energy-efficient, and reliable manner over long-distance networks, several issues such as quality of service, security, metrology, network-resource scheduling and virtualization are being investigated since 15 years. “Computing Networks” explores the core of clusters, grids and clouds networks, giving designers, application developers and users the keys to better construct and use these powerful infrastructures.

This latest textbook from bestselling author, Douglas E. Comer, is a class-tested book providing a comprehensive introduction to cloud computing. Focusing on concepts and principles, rather than commercial offerings by cloud providers and vendors, *The Cloud Computing Book: The Future of Computing Explained* gives readers a complete picture of the advantages and growth of cloud computing, cloud infrastructure, virtualization, automation and orchestration, and cloud-native software design. The book explains real and virtual data center facilities, including computation (e.g., servers, hypervisors, Virtual Machines, and containers), networks (e.g., leaf-spine architecture, VLANs, and VxLAN), and storage mechanisms (e.g., SAN, NAS, and object storage). Chapters on automation and orchestration cover the conceptual organization of systems that automate software deployment and scaling. Chapters on cloud-native software cover parallelism, microservices, MapReduce, controller-based designs, and serverless computing. Although it focuses on concepts and principles, the book uses popular technologies in examples, including Docker containers and Kubernetes. Final chapters explain security in a cloud environment and the use of models to help control the complexity involved in designing software for the cloud. The text is suitable for a one-semester course for software engineers who want to understand cloud, and for IT managers moving an organization’s computing to the cloud.

[Copyright: 4d88a113de1022c94968c817316a7051](#)