

Distributed And Cloud Computing 1st Edition Elsevier

This book constitutes the thoroughly refereed post-conference proceedings of the First International Workshop on Adaptive Resource Management and Scheduling for Cloud Computing, ARMS-CC 2014, held in Conjunction with ACM Symposium on Principles of Distributed Computing, PODC 2014, in Paris, France, in July 2014. The 14 revised full papers (including 2 invited talks) were carefully reviewed and selected from 29 submissions and cover topics such as scheduling methods and algorithms, services and applications, fundamental models for resource management in the cloud.

Cloud Computing erlaubt die Bereitstellung und Nutzung von IT-Infrastruktur, von Plattformen und von Anwendungen aller Art als im Web elektronisch verfügbare Dienste. Die Autoren vermitteln einen Überblick über die Architektur, über Anwendungen und die Entwicklung von Cloud Computing. Ziel ist es, die Leserinnen und Leser auf einen einheitlichen Stand zu bringen und so eine gemeinsame Diskussionsgrundlage zu schaffen. Technische Vorkenntnisse sind für das Verständnis nicht erforderlich.

The Cognitive Approach in Cloud Computing and Internet of Things Technologies for Surveillance Tracking Systems discusses the recent, rapid development of Internet of things (IoT) and its focus on research in smart cities, especially on surveillance tracking systems in which computing devices are widely distributed and huge amounts of dynamic real-time data are collected and processed. Efficient surveillance tracking systems in the Big Data era require

the capability of quickly abstracting useful information from the increasing amounts of data. Real-time information fusion is imperative and part of the challenge to mission critical surveillance tasks for various applications. This book presents all of these concepts, with a goal of creating automated IT systems that are capable of resolving problems without demanding human aid. Examines the current state of surveillance tracking systems, cognitive cloud architecture for resolving critical issues in surveillance tracking systems, and research opportunities in cognitive computing for surveillance tracking systems Discusses topics including cognitive computing architectures and approaches, cognitive computing and neural networks, complex analytics and machine learning, design of a symbiotic agent for recognizing real space in ubiquitous environments, and more Covers supervised regression and classification methods, clustering and dimensionality reduction methods, model development for machine learning applications, intelligent machines and deep learning networks includes coverage of cognitive computing models for scalable environments, privacy and security aspects of surveillance tracking systems, strategies and experiences in cloud architecture and service platform design

This book constitutes the refereed proceedings of the First International Conference on Intelligent, Secure, and Dependable Systems in Distributed and Cloud Environments, ISDDC 2017, held in Vancouver, BC, Canada, in October 2017. The 12 full papers presented together with 1 short paper were carefully reviewed and selected from 43 submissions. This book also contains 3 keynote talks and 2 tutorials. The contributions included in this proceedings cover many aspects of theory and application of effective and efficient paradigms, approaches, and tools for building, maintaining, and managing secure and dependable systems and

infrastructures, such as botnet detection, secure cloud computing and cryptosystems, IoT security, sensor and social network security, behavioral systems and data science, and mobile computing.

Cloud computing offers many advantages to researchers and engineers who need access to high performance computing facilities for solving particular compute-intensive and/or large-scale problems, but whose overall high performance computing (HPC) needs do not justify the acquisition and operation of dedicated HPC facilities. There are, however, a number of fundamental problems which must be addressed, such as the limitations imposed by accessibility, security and communication speed, before these advantages can be exploited to the full. This book presents 14 contributions selected from the International Research Workshop on Advanced High Performance Computing Systems, held in Cetraro, Italy, in June 2012. The papers are arranged in three chapters. Chapter 1 includes five papers on cloud infrastructures, while Chapter 2 discusses cloud applications. The third chapter in the book deals with big data, which is nothing new – large scientific organizations have been collecting large amounts of data for decades – but what is new is that the focus has now broadened to include sectors such as business analytics, financial analyses, Internet service providers, oil and gas, medicine, automotive and a host of others. This book will be of interest to all those whose work involves them with aspects of cloud computing and big data applications.

Distributed and Cloud Computing
From Parallel Processing to the Internet of Things
Morgan Kaufmann

Whether you're already in the cloud, or determining whether or not it makes sense for your organization, *Cloud Computing and Software Services: Theory and Techniques* provides the

technical understanding needed to develop and maintain state-of-the-art cloud computing and software services. From basic concepts and recent research findings to future Big data technologies are used to achieve any type of analytics in a fast and predictable way, thus enabling better human and machine level decision making. Principles of distributed computing are the keys to big data technologies and analytics. The mechanisms related to data storage, data access, data transfer, visualization and predictive modeling using distributed processing in multiple low cost machines are the key considerations that make big data analytics possible within stipulated cost and time practical for consumption by human and machines. However, the current literature available in big data analytics needs a holistic perspective to highlight the relation between big data analytics and distributed processing for ease of understanding and practitioner use. This book fills the literature gap by addressing key aspects of distributed processing in big data analytics. The chapters tackle the essential concepts and patterns of distributed computing widely used in big data analytics. This book also covers the main technologies which support distributed processing. Finally, this book provides insight into applications of big data analytics, highlighting how principles of distributed computing are used in those situations. Practitioners and researchers alike will find this book a valuable tool for their work, helping them to select the appropriate technologies, while understanding the inherent strengths and drawbacks of those technologies. The Encyclopedia of Cloud Computing comprehensively covers all aspects of cloud computing. It provides IT professionals, educators, researchers and students a compendium of cloud computing knowledge – concepts, principles, architecture,

technology, security, privacy and regulatory compliance, applications, adoption, business, and social and legal aspects. Containing contributions from a spectrum of subject matter experts in industry and academia, this unique publication also addresses questions related to technological trends and developments, research opportunities, best practices, standards, and cloud adoption that stakeholders might have in the context of development, operation, management, and use of clouds, providing multiple perspectives. Furthermore, it examines cloud computing's impact now and in the future. The encyclopedia is logically organised into 10 sections and each section into a maximum of 12 chapters, each covering a major topic/area with cross-references as required. The chapters consist of tables, illustrations, side-bars as appropriate. In addition, it also includes highlights at the beginning of each chapter, as well as backend material references and additional resources for further information (including relevant websites, videos and software tools). The encyclopedia also contains illustrations and case studies. A list of acronyms are provided in the beginning and a comprehensive and informative glossary at the end.

All over the world, vast research is in progress on the domain of Industry 4.0 and related techniques. Industry 4.0 is expected to have a very high impact on labor markets, global value chains, education, health, environment, and many social economic aspects. Industry 4.0 Interoperability, Analytics, Security, and Case Studies provides a deeper understanding of the drivers and enablers of Industry 4.0. It includes

real case studies of various applications related to different fields, such as cyber physical systems (CPS), Internet of Things (IoT), cloud computing, machine learning, virtualization, decentralization, blockchain, fog computing, and many other related areas. Also discussed are interoperability, design, and implementation challenges. Researchers, academicians, and those working in industry around the globe will find this book of interest. FEATURES Provides an understanding of the drivers and enablers of Industry 4.0 Includes real case studies of various applications for different fields Discusses technologies such as cyber physical systems (CPS), Internet of Things (IoT), cloud computing, machine learning, virtualization, decentralization, blockchain, fog computing, and many other related areas Covers design, implementation challenges, and interoperability Offers detailed knowledge on Industry 4.0 and its underlying technologies, research challenges, solutions, and case studies

Cloud Computing: Theory and Practice, Second Edition, provides students and IT professionals with an in-depth analysis of the cloud from the ground up. After an introduction to network-centric computing and network-centric content in Chapter One, the book is organized into four sections. Section One reviews basic concepts of concurrency and parallel and distributed systems. Section Two presents such critical components of the cloud ecosystem as cloud service providers, cloud access, cloud data storage, and cloud hardware and software. Section Three covers cloud applications and cloud security, while Section Four presents research topics in cloud

computing. Specific topics covered include resource virtualization, resource management and scheduling, and advanced topics like the impact of scale on efficiency, cloud scheduling subject to deadlines, alternative cloud architectures, and vehicular clouds. An included glossary covers terms grouped in several categories, from general to services, virtualization, desirable attributes and security. Includes new chapters on concurrency, cloud hardware and software, challenges posed by big data and mobile applications and advanced topics Provides a new appendix that presents several cloud computing projects Presents more than 400 references in the text, including recent research results in several areas related to cloud computing ICCCEG 2015, is a main annual research conference aimed at presenting current research being carried out. The idea of the conference is for the scientists, scholars, engineers and students from the Universities all around the world and the industry to present ongoing research activities, and hence to foster research relations between the Universities and the industry.

Developers often struggle when first encountering the cloud. Learning about distributed systems, becoming familiar with technologies such as containers and functions, and knowing how to put everything together can be daunting. With this practical guide, you'll get up to speed on patterns for building cloud native applications and best practices for common tasks such as messaging, eventing, and DevOps. Authors Boris Scholl, Trent Swanson, and Peter Jausovec describe the architectural building blocks

for a modern cloud native application. You'll learn how to use microservices, containers, serverless computing, storage types, portability, and functions. You'll also explore the fundamentals of cloud native applications, including how to design, develop, and operate them. Explore the technologies you need to design a cloud native application Distinguish between containers and functions, and learn when to use them Architect applications for data-related requirements Learn DevOps fundamentals and practices for developing, testing, and operating your applications Use tips, techniques, and best practices for building and managing cloud native applications Understand the costs and trade-offs necessary to make an application portable

Explores key challenges and solutions to assured cloud computing today and provides a provocative look at the face of cloud computing tomorrow This book offers readers a comprehensive suite of solutions for resolving many of the key challenges to achieving high levels of assurance in cloud computing. The distillation of critical research findings generated by the Assured Cloud Computing Center of Excellence (ACC-UCoE) of the University of Illinois, Urbana-Champaign, it provides unique insights into the current and future shape of robust, dependable, and secure cloud-based computing and data cyberinfrastructures. A survivable and distributed cloud-computing-based infrastructure can enable the configuration of any dynamic systems-of-systems that contain both trusted and partially trusted resources and services sourced from multiple organizations. To assure mission-critical computations and workflows that rely on such

systems-of-systems it is necessary to ensure that a given configuration does not violate any security or reliability requirements. Furthermore, it is necessary to model the trustworthiness of a workflow or computation fulfillment to a high level of assurance. In presenting the substance of the work done by the ACC-UCoE, this book provides a vision for assured cloud computing illustrating how individual research contributions relate to each other and to the big picture of assured cloud computing. In addition, the book: Explores dominant themes in cloud-based systems, including design correctness, support for big data and analytics, monitoring and detection, network considerations, and performance Synthesizes heavily cited earlier work on topics such as DARE, trust mechanisms, and elastic graphs, as well as newer research findings on topics, including R-Storm, and RAMP transactions Addresses assured cloud computing concerns such as game theory, stream processing, storage, algorithms, workflow, scheduling, access control, formal analysis of safety, and streaming Bringing together the freshest thinking and applications in one of today's most important topics, Assured Cloud Computing is a must-read for researchers and professionals in the fields of computer science and engineering, especially those working within industrial, military, and governmental contexts. It is also a valuable reference for advanced students of computer science.

Road accidents caused by impaired and distracted driving as well as traffic congestion are on the rise, with the numbers increasing dramatically every day. Intelligent

transportation systems (ITS) aim to improve the efficiency and safety of traveling by consolidating vehicle operations, managing vehicle traffic, and notifying drivers with alerts and safety messages in real time. Vehicular Cloud Computing for Traffic Management and Systems provides innovative research on the rapidly advancing applications of vehicle-to-vehicle and vehicle-to-infrastructure communication. It also covers the need to fully utilize vehicular ad-hoc network (VANET) resources to provide updated and dynamic information about the conditions of road traffic so that the number of road accidents can be minimized. Featuring research on topics such as identity management, computational architecture, and resource management, this book is ideally designed for urban planners, researchers, policy makers, graduate-level students, transportation engineers, and technology developers seeking current research on vehicle computational design, architecture, security, and privacy. 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.

This book constitutes the proceedings of the 23rd International Conference on Parallel and Distributed Computing, Euro-Par 2017, held in Santiago de Compostela, Spain, in August/September 2017. The 50 revised full papers presented together with 2 abstract of invited talks and 1 invited paper were carefully reviewed and selected from 176 submissions. The papers are organized in the following topical sections: support tools and environments; performance and power modeling, prediction and evaluation; scheduling and load balancing; high performance architectures and compilers; parallel and distributed data management and analytics; cluster and cloud computing; distributed systems and algorithms; parallel and distributed programming, interfaces and languages; multicore and manycore parallelism; theory and algorithms for parallel computation and networking; parallel numerical methods and applications; and accelerator computing.

Great POSSIBILITIES and high future prospects to become ten times folds in the near FUTURE DESCRIPTION The book “Handbook of Cloud Computing” provides the latest and in-depth information of this relatively new and another platform for scientific computing which has great possibilities and high future prospects to become ten folds in near future. The book covers in comprehensive manner all aspects and terminologies associated with cloud computing like SaaS, PaaS and IaaS and also elaborates almost every cloud computing service model. The book highlights several other aspects of cloud computing like Security, Resource allocation, Simulation Platforms and futuristic trend i.e. Mobile cloud computing. The book will benefit all the readers with all in-depth technical information which is required to understand current and futuristic concepts of cloud computing. No prior knowledge of cloud computing or any of its related technology is required in reading this book. KEY FEATURES Comprehensively gives clear picture of current state-of-the-art aspect of cloud computing by elaborating terminologies, models and other related terms. Enlightens all major players in Cloud Computing industry providing services in terms of SaaS, PaaS and IaaS. Highlights Cloud Computing Simulators, Security Aspect and Resource Allocation. In-depth presentation with well-illustrated diagrams and simple to understand technical concepts of cloud. WHAT WILL YOU LEARN Cloud Computing, Virtualisation Software as a Service, Platform as a Service, Infrastructure as a Service Data in Cloud and its Security Cloud Computing – Simulation, Mobile Cloud Computing Specific Cloud Service Models Resource Allocation in Cloud Computing WHO THIS BOOK IS FOR Students of Polytechnic Diploma Classes- Computer Science/ Information Technology Graduate Students- Computer Science/ CSE / IT/ Computer Applications Master Class Students—Msc (CS/IT)/ MCA/ M.Phil, M.Tech, M.S.

Researcher's—Ph.D Research Scholars doing work in Virtualization, Cloud Computing and Cloud Security Industry Professionals- Preparing for Certifications, Implementing Cloud Computing and even working on Cloud Security Table of Contents 1. Introduction to Cloud Computing 2. Virtualisation 3. Software as a Service 4. Platform as a Service 5. Infrastructure as a Service 6. Data in Cloud 7. Cloud Security 8. Cloud Computing – Simulation 9. Specific Cloud Service Models 10. Resource Allocation in Cloud Computing 11. Mobile Cloud Computing

This book addresses topics related to cloud and Big Data technologies, architecture and applications including distributed computing and data centers, cloud infrastructure and security, and end-user services. The majority of the book is devoted to the security aspects of cloud computing and Big Data. Cloud computing, which can be seen as any subscription-based or pay-per-use service that extends the Internet's existing capabilities, has gained considerable attention from both academia and the IT industry as a new infrastructure requiring smaller investments in hardware platforms, staff training, or licensing software tools. It is a new paradigm that has ushered in a revolution in both data storage and computation. In parallel to this progress, Big Data technologies, which rely heavily on cloud computing platforms for both data storage and processing, have been developed and deployed at breathtaking speed. They are among the most frequently used technologies for developing applications and services in many fields, such as the web, health, and energy. Accordingly, cloud computing and Big Data technologies are two of the most central current and future research mainstreams. They involve and impact a host of fields, including business, scientific research, and public and private administration. Gathering extended versions of the best papers presented at the Third

International Conference on Cloud Computing Technologies and Applications (CloudTech'17), this book offers a valuable resource for all Information System managers, researchers, students, developers, and policymakers involved in the technological and application aspects of cloud computing and Big Data.

This SpringerBrief reviews the existing market-oriented strategies for economically managing resource allocation in distributed systems. It describes three new schemes that address cost-efficiency, user incentives, and allocation fairness with regard to different scheduling contexts. The first scheme, taking the Amazon EC2™ market as a case of study, investigates the optimal resource rental planning models based on linear integer programming and stochastic optimization techniques. This model is useful to explore the interaction between the cloud infrastructure provider and the cloud resource customers. The second scheme targets a free-trade resource market, studying the interactions amongst multiple rational resource traders. Leveraging an optimization framework from AI, this scheme examines the spontaneous exchange of resources among multiple resource owners. Finally, the third scheme describes an experimental market-oriented resource sharing platform inspired by eBay's transaction model. The study presented in this book sheds light on economic models and their implication to the utility-oriented scheduling problems.

Cloud computing has recently emerged as one of the buzzwords in the ICT industry. Numerous IT vendors are promising to offer computation, storage, and application hosting services and to provide coverage in several continents, offering service-level agreements (SLA)-backed performance and uptime promises for their services. While these "clouds" are the natural evolution of traditional data centers, they are distinguished by exposing resources

(computation, data/storage, and applications) as standards-based Web services and following a "utility" pricing model where customers are charged based on their utilization of computational resources, storage, and transfer of data. This book explains the importance and fundamentals of Cloud Computing Concepts

Distributed systems intertwine with our everyday lives. The benefits and current shortcomings of the underpinning technologies are experienced by a wide range of people and their smart devices. With the rise of large-scale IoT and similar distributed systems, cloud bursting technologies, and partial outsourcing solutions, private entities are encouraged to increase their efficiency and offer unparalleled availability and reliability to their users. The Research Anthology on Architectures, Frameworks, and Integration Strategies for Distributed and Cloud Computing is a vital reference source that provides valuable insight into current and emergent research occurring within the field of distributed computing. It also presents architectures and service frameworks to achieve highly integrated distributed systems and solutions to integration and efficient management challenges faced by current and future distributed systems. Highlighting a range of topics such as data sharing, wireless sensor networks, and scalability, this multi-volume book is ideally designed for system administrators, integrators, designers, developers, researchers, academicians, and students.

This practically-focused reference presents a comprehensive overview of the state of the art in Cloud Computing, and examines the potential for future Cloud and Cloud-related technologies to address specific industrial and research challenges. This new edition explores both established and emergent principles, techniques, protocols and algorithms involved with the design, development, and management of Cloud-based systems. The text reviews a range of

applications and methods for linking Clouds, undertaking data management and scientific data analysis, and addressing requirements both of data analysis and of management of large scale and complex systems. This new edition also extends into the emergent next generation of mobile telecommunications, relating network function virtualization and mobile edge Cloud Computing, as supports Smart Grids and Smart Cities. As with the first edition, emphasis is placed on the four quality-of-service cornerstones of efficiency, scalability, robustness, and security.

Clouds are distributed technology platforms that leverage sophisticated technology innovations to provide highly scalable and resilient environments that can be remotely utilized by organizations in a multitude of powerful ways. To successfully build upon, integrate with, or even create a cloud environment requires an understanding of its common inner mechanics, architectural layers, and models, as well as an understanding of the business and economic factors that result from the adoption and real-world use of cloud-based services. In *Cloud Computing: Concepts, Technology & Architecture*, Thomas Erl, one of the world's top-selling IT authors, teams up with cloud computing experts and researchers to break down proven and mature cloud computing technologies and practices into a series of well-defined concepts, models, technology mechanisms, and technology architectures, all from an industry-centric and vendor-neutral point of view. In doing so, the book establishes concrete, academic coverage with a focus on structure, clarity, and well-defined building blocks for mainstream cloud computing platforms and solutions. Subsequent to technology-centric coverage, the book proceeds to establish business-centric models and metrics that allow for the financial assessment of cloud-based IT resources and their comparison to those hosted on traditional IT

enterprise premises. Also provided are templates and formulas for calculating SLA-related quality-of-service values and numerous explorations of the SaaS, PaaS, and IaaS delivery models. With more than 260 figures, 29 architectural models, and 20 mechanisms, this indispensable guide provides a comprehensive education of cloud computing essentials that will never leave your side.

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 e-commerce 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 first textbook to teach students how to build data analytic solutions on large data sets using cloud-based technologies. This is the first textbook to teach students how to build data analytic solutions on large data sets (specifically in Internet of Things applications) using cloud-based technologies for data storage, transmission and mashup, and AI techniques to analyze this data. This textbook is designed to train college students to master modern cloud computing systems in operating principles, architecture design, machine learning algorithms, programming models and software tools for big data mining, analytics, and cognitive applications. The book will be suitable for use in one-semester computer science or electrical engineering courses on cloud computing, machine learning, cloud programming, cognitive computing, or big data science. The book will also be very useful as a reference for professionals who want to work in cloud computing and data science. Cloud and Cognitive Computing begins with two introductory chapters on fundamentals of cloud computing, data science, and adaptive computing that lay the foundation for the rest of the book. Subsequent chapters cover topics including cloud architecture, mashup services, virtual machines, Docker containers, mobile clouds, IoT and AI, inter-cloud mashups, and cloud performance and benchmarks, with a focus on Google's Brain Project, DeepMind, and X-Lab programs, IBKai HwangM SyNapse, Bluemix programs, cognitive initiatives, and neurocomputers. The book

then covers machine learning algorithms and cloud programming software tools and application development, applying the tools in machine learning, social media, deep learning, and cognitive applications. All cloud systems are illustrated with big data and cognitive application examples.

This book constitutes the refereed proceedings of the First International Conference on Advances in Parallel, Distributed Computing Technologies and Applications, PDCTA 2011, held in Tirunelveli, India, in September 2011. The 64 revised full papers were carefully reviewed and selected from over 400 submissions. Providing an excellent international forum for sharing knowledge and results in theory, methodology and applications of parallel, distributed computing the papers address all current issues in this field with special focus on algorithms and applications, computer networks, cyber trust and security, wireless networks, as well as mobile computing and bioinformatics.

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, - lowing 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.

Provides an up-to-date analysis of big data and multi-agent systems The term Big Data refers to the cases, where data sets are too large or too complex for traditional data-processing software. With the spread of new concepts such as Edge Computing or the Internet of Things, production, processing and consumption of this data becomes more and more distributed. As a result, applications increasingly require multiple agents that can work together. A multi-agent system (MAS) is a self-organized computer system that comprises multiple intelligent agents interacting to solve problems that are beyond the capacities of individual agents. Modern Big Data Architectures examines modern concepts and architecture for Big Data processing and analytics. This unique, up-to-date volume provides joint analysis of big data and multi-agent systems, with emphasis on distributed, intelligent processing of very large data sets. Each chapter contains practical examples and detailed solutions suitable for a wide variety of applications. The author, an internationally-recognized expert in Big Data and distributed Artificial Intelligence, demonstrates how base concepts such as agent, actor, and micro-service have reached a point of convergence—enabling next generation systems to be built by incorporating the best aspects of the field. This book: Illustrates how data sets are produced

and how they can be utilized in various areas of industry and science Explains how to apply common computational models and state-of-the-art architectures to process Big Data tasks Discusses current and emerging Big Data applications of Artificial Intelligence Modern Big Data Architectures: A Multi-Agent Systems Perspective is a timely and important resource for data science professionals and students involved in Big Data analytics, and machine and artificial learning.

A highly accessible reference offering a broad range of topics and insights on large scale network-centric distributed systems Evolving from the fields of high-performance computing and networking, large scale network-centric distributed systems continues to grow as one of the most important topics in computing and communication and many interdisciplinary areas. Dealing with both wired and wireless networks, this book focuses on the design and performance issues of such systems. Large Scale Network-Centric Distributed Systems provides in-depth coverage ranging from ground-level hardware issues (such as buffer organization, router delay, and flow control) to the high-level issues immediately concerning application or system users (including parallel programming, middleware, and OS support for such computing systems). Arranged in five parts, it explains and analyzes complex topics to an unprecedented degree: Part 1: Multicore and Many-Core (Mc) Systems-on-Chip Part 2: Pervasive/Ubiquitous Computing and Peer-to-Peer Systems Part 3: Wireless/Mobile Networks Part 4: Grid and Cloud Computing Part 5: Other Topics Related to Network-Centric Computing and Its Applications Large Scale Network-Centric Distributed Systems is an incredibly useful resource for practitioners, postgraduate students, postdocs, and researchers.

This book reviews the challenging issues that present barriers to greater implementation of the

cloud computing paradigm, together with the latest research into developing potential solutions. Topics and features: presents a focus on the most important issues and limitations of cloud computing, covering cloud security and architecture, QoS and SLAs; discusses a methodology for cloud security management, and proposes a framework for secure data storage and identity management in the cloud; introduces a simulation tool for energy-aware cloud environments, and an efficient congestion control system for data center networks; examines the issues of energy-aware VM consolidation in the IaaS provision, and software-defined networking for cloud related applications; reviews current trends and suggests future developments in virtualization, cloud security, QoS data warehouses, cloud federation approaches, and DBaaS provision; predicts how the next generation of utility computing infrastructures will be designed.

Innovations in cloud and service-oriented architectures continue to attract attention by offering interesting opportunities for research in scientific communities. Although advancements such as computational power, storage, networking, and infrastructure have aided in making major progress in the implementation and realization of cloud-based systems, there are still significant concerns that need to be taken into account. Principles, Methodologies, and Service-Oriented Approaches for Cloud Computing aims to present insight into Cloud principles, examine associated methods and technologies, and investigate the use of service-oriented computing technologies. In addressing supporting infrastructure of the Cloud, including associated challenges and pressing issues, this reference source aims to present researchers, engineers, and IT professionals with various approaches in Cloud computing.

Technology trends may come and go, but cloud computing technologies have been gaining

consideration in the commercial world due to its ability to provide on-demand access to resources, control the software environment, and supplement existing systems. *Pervasive Cloud Computing Technologies: Future Outlooks and Interdisciplinary Perspectives* explores the latest innovations with cloud computing and the impact of these new models and technologies. This book will present case studies and research on the future of cloud computing technologies and its ability to increase connectivity of various entities of the world. It is an essential resource for technology practitioners, engineers, managers, and academics aiming to gain the knowledge of these novel and pervasive technologies.

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. 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.

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 is a buzz-word in today's information technology (IT) that nobody can escape. But what is really behind it? There are many interpretations of this term, but no standardized or even uniform definition. Instead, as a result of the multi-faceted viewpoints and the diverse interests expressed by the various stakeholders, cloud computing is perceived as a rather fuzzy concept. With this book, the authors deliver an overview of cloud computing architecture, services, and applications. Their aim is to bring readers up to date on this technology and thus to provide a common basis for discussion, new research, and novel application scenarios. They first introduce the foundation of cloud computing with its basic technologies, such as virtualization and Web services. After that they discuss the cloud architecture and its service modules. The following chapters then cover selected commercial cloud offerings (including Amazon Web Services and Google App Engine) and management tools, and present current related open-source developments (including Hadoop, Eucalyptus, and Open Cirrus™). Next, economic considerations (cost and business models) are discussed, and an evaluation of the cloud market situation is given. Finally, the appendix contains some practical examples of how to use cloud resources or cloud applications, and a glossary provides concise definitions of key terms. The authors' presentation does not require in-depth technical knowledge. It is equally intended as an introduction for students in software engineering, web technologies, or business development, for professional software developers or system architects, and for future-oriented decision-makers like top executives and managers.

This book constitutes the thoroughly refereed post-conference proceedings of the First

International Workshop on Algorithmic Aspects of Cloud Computing, ALGO CLOUD 2015, held in Patras, Greece, in September 2015 in conjunction with ALGO 2015. The 13 revised full papers presented together with 2 tutorial papers were carefully reviewed and selected from 37 initial submissions. They cover a wide range of topics in two main tracks: algorithmic aspects of large-scale data stores, and software tools and distributed architectures for cloud-based data management.

This book constitutes the thoroughly refereed proceedings of the 5th International Conference on Cloud Computing and Services Science, CLOSER 2015, held in Lisbon, Portugal, in May 2015. The 14 revised full papers presented together with one invited paper were selected from 146 paper submissions. The papers focus on the following topics: cloud computing fundamentals; services science foundations for cloud computing; cloud computing platforms and applications; cloud computing enabling technologies; and mobile cloud computing services.

[Copyright: 4d88a113de1022c94968c817316a7051](https://doi.org/10.1007/978-3-319-21131-6)