

2011 IEEE Third International Conference on Cloud Computing Technology and Science

(CloudCom 2011)

Athens, Greece

29 November – 1 December 2011



IEEE Catalog Number: CFP11CLU-PRT
ISBN: 978-1-4673-0090-2

2011 Third IEEE International Conference on Cloud Computing Technology and Science

CloudCom 2011

Table of Contents

Message from the General Chairs.....	xiv
Message from the Program Committee Chairs.....	xv
Message from the Cloud Computing Association - CloudCom.org.....	xvi
Message from the Workshop Chair.....	xvii
Conference Committees.....	xviii
External Reviewers.....	xxiv
Workshops Committees.....	xxvi

eScience

Can Cloud Computing Be Used for Planning? An Initial Study	1
<i>Qiang Lu, You Xu, Ruoyun Huang, Yixin Chen, and Guoliang Chen</i>	
Ordinal Optimized Scheduling of Scientific Workflows in Elastic Compute Clouds	9
<i>Fan Zhang, Junwei Cao, Kai Hwang, and Cheng Wu</i>	
Scientific Workflow Makespan Reduction through Cloud Augmented Desktop Grids	18
<i>Christopher J. Reynolds, Stephen Winter, Gabor Z. Terstyanszky, Tamas Kiss, Pamela Greenwell, Sandor Acs, and Peter Kacsuk</i>	
Computational Neuroscience as a Service: Porting MIIND to the Cloud	24
<i>Björn-Ole Gerckens, Karim Djemame, and Marc de Kamps</i>	

MapReduce

Byzantine Fault-Tolerant MapReduce: Faults are Not Just Crashes	32
<i>Pedro Costa, Marcelo Pasin, Alysson N. Bessani, and Miguel Correia</i>	
Matchmaking: A New MapReduce Scheduling Technique	40
<i>Chen He, Ying Lu, and David Swanson</i>	
Scalable and Low-Latency Data Processing with Stream MapReduce	48
<i>Andrey Brito, André Martin, Thomas Knauth, Stephan Creutz, Diogo Becker, Stefan Weigert, and Christof Fetzer</i>	
Optimizing Multiple Machine Learning Jobs on MapReduce	59
<i>Hiroshi Tamano, Shinji Nakadai, and Takuya Araki</i>	

Applications and Evaluation of In-memory MapReduce	67
<i>Kim-Thomas Rehmman and Michael Schöettner</i>	
Parallel Outlier Detection Using KD-Tree Based on MapReduce	75
<i>Qing He, Yunlong Ma, Qun Wang, Fuzhen Zhuang, and Zhongzhi Shi</i>	
iHadoop: Asynchronous Iterations for MapReduce	81
<i>Eslam Elnikety, Tamer Elsayed, and Hany E. Ramadan</i>	
Virtualization	
Comparing VM-Placement Algorithms for On-Demand Clouds	91
<i>K. Mills, J. Filliben, and C. Dabrowski</i>	
Optimization-Based Virtual Machine Manager for Private Cloud Computing	99
<i>Dusit Niyato</i>	
Combinatorial Auction-Based Dynamic VM Provisioning and Allocation in Clouds	107
<i>Sharrukh Zaman and Daniel Grosu</i>	
Energy-Aware Task Consolidation Technique for Cloud Computing	115
<i>Ching-Hsien Hsu, Shih-Chang Chen, Chih-Chun Lee, Hsi-Ya Chang, Kuan-Chou Lai, Kuan-Ching Li, and Chunming Rong</i>	
On Improvement of Cloud Virtual Machine Availability with Virtualization Fault Tolerance Mechanism	122
<i>Chao-Tung Yang, Wei-Li Chou, Ching-Hsien Hsu, and Alfredo Cuzzocrea</i>	
Applications Know Best: Performance-Driven Memory Overcommit with Ginkgo	130
<i>Michael R. Hines, Abel Gordon, Marcio Silva, Dilma Da Silva, Kyung Ryu, and Muli Ben-Yehuda</i>	
Autonomic Resource Management Handling Delayed Configuration Effects	138
<i>Oliver Niehöerster and Andre Brinkmann</i>	
Characterizing Power and Energy Usage in Cloud Computing Systems	146
<i>Ziming Zhang and Song Fu</i>	
Efficiently Synchronizing Virtual Machines in Cloud Computing Environments	154
<i>Shuntaro Tonosaki, Hiroshi Yamada, and Kenji Kono</i>	
Modeling for Dynamic Cloud Scheduling Via Migration of Virtual Machines	163
<i>Wubin Li, Johan Tordsson, and Erik Elmroth</i>	
Security and Privacy	
A Cryptographic Protocol for Communication in a Redundant Array of Independent Net-storages	172
<i>Martin Gilje Jaatun, Gansen Zhao, and Stian Alapnes</i>	
A Multi-Level Security Model for Partitioning Workflows over Federated Clouds	180
<i>Paul Watson</i>	
As Strong as the Weakest Link: Handling Compromised Components in OpenStack	189
<i>Aryan Taheri Monfared and Martin Gilje Jaatun</i>	

An Autonomous Agent Based Incident Detection System for Cloud Environments	197
<i>Frank Doelitzscher, Christoph Reich, Martin Knahl, and Nathan Clarke</i>	
Implementation of a Secure Genome Sequence Search Platform on Public Cloud: Leveraging Open Source Solutions	205
<i>Shyam Kumar Doddavula and Vikas Saxena</i>	
Obtaining the Integrity of Your Virtual Machine in the Cloud	213
<i>Aimin Yu, Yu Qin, and Dan Wang</i>	
Privacy-preserving Collaborative Filtering for the Cloud	223
<i>Anirban Basu, Jaideep Vaidya, Hiroaki Kikuchi, and Theo Dimitrakos</i>	
A Quantitative Analysis of Current Security Concerns and Solutions for Cloud Computing	231
<i>Nelson Gonzalez, Charles Miers, Fernando Redígolo, Tereza Carvalho, Marcos Simplicio, Mats Näslund, and Makan Pourzandi</i>	
Checking Running and Dormant Virtual Machines for the Necessity of Security Updates in Cloud Environments	239
<i>Roland Schwarzkopf, Matthias Schmidt, Christian Strack, and Bernd Freisleben</i>	
myTrustedCloud: Trusted Cloud Infrastructure for Security-critical Computation and Data Managment	247
<i>David Wallom, Matteo Turilli, Andrew Martin, Anbang Raun, Gareth Taylor, Nigel Hargreaves, and Alan McMoran</i>	
Security Infrastructure for On-demand Provisioned Cloud Infrastructure Services	255
<i>Yuri Demchenko, Canh Ngo, Cees de Laat, Tomasz Wiktor Wlodarczyk, Chunming Rong, and Wolfgang Ziegler</i>	
Multi-User Private Keyword Search for Cloud Computing	264
<i>Yanjiang Yang, Haibing Lu, and Jian Weng</i>	
Interoperability and Standardization	
A Rule-based Approach for Infrastructure Providers' Interoperability	272
<i>Jorge Ejarque, Javier Álvarez, Raül Sirvent, and Rosa M. Badia</i>	
A Semantic Interoperability Framework for Cloud Platform as a Service	280
<i>Nikolaos Loutas, Eleni Kamateri, and Konstantinos Tarabanis</i>	
Services and Applications	
QoS-aware SLA-based Advanced Reservation of Infrastructure as a Service	288
<i>Kuan Lu, Thomas Röblitz, Ramin Yahyapour, Edwin Yaqub, and Constantinos Kotsokalis</i>	
QoS-control of Structured Parallel Computations: A Predictive Control Approach	296
<i>Gabriele Mencagli and Marco Vanneschi</i>	
Deconstructing Amazon EC2 Spot Instance Pricing	304
<i>Orna Agmon Ben-Yehuda, Muli Ben-Yehuda, Assaf Schuster, and Dan Tsafir</i>	

A Framework for Dynamic Resource Provisioning and Adaptation in IaaS Clouds	312
<i>Ta Nguyen Binh Duong, Xiaorong Li, and Rick Siow Mong Goh</i>	
Cost-Efficient Scheduling Heuristics for Deadline Constrained Workloads on Hybrid Clouds	320
<i>Ruben Van den Bossche, Kurt Vanmechelen, and Jan Broeckhove</i>	
Towards a Contextualization Solution for Cloud Platform Services	328
<i>Django Armstrong, Karim Djemame, Srijith Nair, Johan Tordsson, and Wolfgang Ziegler</i>	
Implementation of a Fast Vector Packing Algorithm and its Application for Server Consolidation	332
<i>Shyam Kumar Doddavula, Mudit Kaushik, and Akansha Jain</i>	
A Cloud Infrastructure for Collaborative Digital Public Services	340
<i>Yannis Charalabidis, Sotirios Koussouris, and Antonis Ramfos</i>	
Scalable Service Containers	348
<i>Sami Yangui, Mohamed Mohamed, Samir Tata, and Samir Moalla</i>	
A Cloud Environment for Data-intensive Storage Services	357
<i>Elliot K. Kolodner, Sivan Tal, Dimosthenis Kyriazis, Dalit Naor, Miriam Allalouf, Lucia Bonelli, Per Brand, Albert Eckert, Erik Elmroth, Spyridon V. Gogouvitis, Danny Harnik, Francisco Hernandez, Michael C. Jaeger, Ewnetu Bayuh Lakew, Jose Manuel Lopez, Mirko Lorenz, Alberto Messina, Alexandra Shulman-Peleg, Roman Talyansky, Athanasios Voulodimos, and Yaron Wolfsthal</i>	
A Cloud Framework for Parameter Sweeping Data Mining Applications	367
<i>Fabrizio Marozzo, Domenico Talia, and Paolo Trunfio</i>	
A Cloud-unaware Programming Model for Easy Development of Composite Services	375
<i>Enric Tejedor, Jorge Ejarque, Francesc Lordan, Roger Rafanell, Javier Álvarez, Daniele Lezzi, Raül Sirvent, and Rosa M. Badia</i>	
Converting a High Performance Application to an Elastic Cloud Application	383
<i>Dinesh Rajan, Anthony Canino, Jesus A. Izaguirre, and Douglas Thain</i>	
Architecture	
EQS: An Elastic and Scalable Message Queue for the Cloud	391
<i>Nam-Luc Tran, Sabri Skhiri, and Esteban Zimányi</i>	
CloudML: An Integrated Language for Resource, Service and Request Description for D-Clouds	399
<i>Glauco Gonçalves, Patricia Endo, Marcelo Santos, Djamel Sadok, Judith Kelner, Bob Melander, and Jan-Erik Mångs</i>	
A-Tree: Distributed Indexing of Multidimensional Data for Cloud Computing Environments	407
<i>Andreas Papadopoulos and Dimitrios Katsaros</i>	
QoS-aware Deployment of Network of Virtual Appliances Across Multiple Clouds	415
<i>Amir Vahid Dastjerdi, Saurabh Kumar Garg, and Rajkumar Buyya</i>	

Energy-Efficient Resource Management for Cloud Computing Infrastructures	424
<i>Marco Guazzone, Cosimo Anglano, and Massimo Canonico</i>	

Fusion of Public and Private Research and Innovation

A User-Oriented, Customizable Infrastructure Sharing Approach for Hybrid Cloud Computing Environments	432
<i>Prodromos Makris, Dimitrios N. Skoutas, Panagiotis Rizomiliotis, and Charalabos Skianis</i>	
Managing Wearable Sensor Data through Cloud Computing	440
<i>Charalampos Doukas and Ilias Maglogiannis</i>	

General

How to Track Your Data: The Case for Cloud Computing Provenance	446
<i>Olive Qing Zhang, Markus Kirchberg, Ryan K.L. Ko, and Bu Sung Lee</i>	
Taxonomy and Requirements Rationalization for Infrastructure in Cloud-based Software Testing	454
<i>Philip Robinson and Carmelo Ragusa</i>	
A Business Resolution Engine for Cloud Marketplaces	462
<i>Andreas Menychtas, Anna Gatzoura, and Theodora Varvarigou</i>	

Short Papers

Research in Progress and Applications

Impact of Resource over-Reservation (ROR) and Dropping Policies on Cloud Resource Allocation	470
<i>Felipe Díaz, Elias A. Doumith, and Maurice Gagnaire</i>	
Towards Economic Energy Trading in Cloud Environments	477
<i>Andreas Zinnen and Thomas Engel</i>	
Scoring System Utilization through Business Profiles	482
<i>Jesus Omana Iglesias, James Thorburn, Trevor Parsons, John Murphy, and Patrick O'Sullivan</i>	
A Conceptual Framework of Deploying Cloud IaaS in Higher Educational Institutions	489
<i>Mary-Jane Sule</i>	
Combining Grid and Cloud Resources for Hybrid Scientific Computing Executions	494
<i>Amanda Calatrava, Germán Moltó, and Vicente Hernández</i>	
Autonomic Cloud Computing: Giving Intelligence to Simpleton Nodes	502
<i>Patricia Takako Endo, Djamel Sadok, and Judith Kelter</i>	
A Deployment Platform for Dynamically Scaling Applications in the Cloud	506
<i>Rui Han, Li Guo, Yike Guo, and Sijin He</i>	
A Front-end, Hadoop-based Data Management Service for Efficient Federated Clouds	511
<i>George Kousiouris, George Vafiadis, and Theodora Varvarigou</i>	

Infrastructure Deployment Over the Cloud	517
<i>Carlos de Alfonso, Miguel Caballer, Fernando Alvarruiz, Germán Moltó, and Vicente Hernández</i>	
Mapping of Cloud Standards to the Taxonomy of Interoperability in IaaS	522
<i>Ralf Teckelmann, Christoph Reich, and Anthony Sulistio</i>	
Managing Incidents in Smart Grids à la Cloud	527
<i>Cristina Alcaraz, Isaac Agudo, David Nuñez, and Javier Lopez</i>	
An Integrated Cloud-based Healthcare Infrastructure	532
<i>E. Ekonomou, L. Fan, W. Buchanan, and C. Thüemmler</i>	
DIRAC Integration with Cloud Stack	537
<i>V. Fernandez Albor, J. Saborido, F. Gomez-Folgar, J. López Cacheiro, and R. Graciani Díaz</i>	
High Performance Live Migration through Dynamic Page Transfer Reordering and Compression	542
<i>Petter Svärd, Johan Tordsson, Benoit Hudzia, and Erik Elmroth</i>	
Efficient Migration of Virtual Machines between Public and Private Cloud	549
<i>Chun-Hui Suen, Markus Kirchberg, and Bu Sung Lee</i>	
VM Leakage and Orphan Control in Open-Source Clouds	554
<i>C. Dabrowski and K. Mills</i>	
FutureGrid Image Repository: A Generic Catalog and Storage System for Heterogeneous Virtual Machine Images	560
<i>Javier Diaz, Gregor von Laszewski, Fugang Wang, Andrew J. Younge, and Geoffrey Fox</i>	
Using Mahout for Clustering Wikipedia's Latest Articles: A Comparison between K-means and Fuzzy C-means in the Cloud	565
<i>Rui Máximo Esteves and Chunming Rong</i>	
Locality-Aware Reduce Task Scheduling for MapReduce	570
<i>Mohammad Hammoud and Majd F. Sakr</i>	
Drug Repositioning on the Cloud	577
<i>Vassilis Virvilis, Anuj Sharma, and Miltiadis Koutsokeras</i>	
Towards Reproducible eScience in the Cloud	582
<i>Jonathan Klinginsmith, Malika Mahoui, and Yuqing Melanie Wu</i>	
Goal-oriented Self-management of In-memory Distributed Data Grid Platforms	587
<i>Liliana Rosa, Luís Rodrigues, and Antónia Lopes</i>	
Dynamic Resource Allocation in Cloud Environment Under Time-variant Job Requests	592
<i>Davide Tamaro, Elias A. Doumith, Sawsan Al Zahr, Jean-Paul Smets, and Maurice Gagnaire</i>	
A Bi-Criteria Truthful Mechanism for Scheduling of Workflows in Clouds	599
<i>Hamid Mohammadi Fard, Radu Prodan, Georg Moser, and Thomas Fahringer</i>	

Workshop Papers

HPCCLOUD

Leveraging on High-Performance Computing and Cloud Technologies in Digital Libraries: A Case Study	606
<i>Peter Wittek and Sándor Darányi</i>	
An Initial Survey on Integration and Application of Cloud Computing to High Performance Computing	612
<i>Tomasz Wiktor Włodarczyk and Chunming Rong</i>	
A Framework for Cloud-Based Large-Scale Data Analytics and Visualization: Case Study on Multiscale Climate Data	618
<i>Sifei Lu, Reuben Mingguang Li, William Chandra Tjhi, Kee Khoon Lee, Long Wang, Xiaorong Li, and Di Ma</i>	
Visual Modeling for Parallel Programming Based on DSL	623
<i>Wang Mengmeng, Yu Ce, Sun Jizhou, Sun Chao, and Chen Jinyan</i>	
Integer Factorization Using Hadoop	628
<i>Son T. Nguyen, Semere Tsehay Ghebreorgish, Nour Alabbasi, and Chunming Rong</i>	
Fast n-point Correlation Function Approximation with Recursive Convolution for Scalar Fields	634
<i>Xiang Zhang and Ce Yu</i>	
Towards Multi-Service Traffic Shaping in Two-Tier Enterprise Data Centers	640
<i>Yesid Jarma, Marcelo Dias de Amorim, and Yannis Viniotis</i>	
Energy Efficient Free Cooling System for Data Centers	646
<i>Christy Sujatha D. and Satheesh Abimannan</i>	
Viteraas: Virtual Cluster as a Service	652
<i>Frank Doelitzscher, Markus Held, Christoph Reich, and Anthony Sulistio</i>	
Automating Application Deployment in Infrastructure Clouds	658
<i>Gideon Juve and Ewa Deelman</i>	

NETCLOUD

GEMBus as a Service Oriented Platform for Cloud-Based Composable Services	666
<i>Mary Grammatikou, Constantinos Marinos, Yuri Demchenko, Diego R. Lopez, Krzysztof Dombek, and Jordi Jofre</i>	
An OpenFlow Based Network Virtualization Framework for the Cloud	672
<i>Jon Matias, Eduardo Jacob, David Sanchez, and Yuri Demchenko</i>	
A Network Virtualization Framework for IP Infrastructure Provisioning	679
<i>Bo Peng, Ali Hammad, Reza Nejabati, Siamak Azodolmolky, Dimitra Simeonidou, and Victor Reijs</i>	
Virtual Infrastructure Planning in Elastic Cloud Deploying Optical Networking	685
<i>Markos P. Anastasopoulos, Anna Tzanakaki, and Konstantinos Georgakilas</i>	

Provisioning and Evaluating Multi-domain Networked Clouds for Hadoop-based Applications	690
<i>Anirban Mandal, Yufeng Xin, Ilia Baldine, Paul Ruth, Chris Heerman, Jeff Chase, Victor Orlikowski, and Aydan Yumerefendi</i>	
Security Framework for Virtualised Infrastructure Services Provisioned On-demand	698
<i>Canh Ngo, Peter Membrey, Yuri Demchenko, and Cees de Laat</i>	
GridARS: A Grid Advanced Resource Management System Framework for Intercloud	705
<i>Atsuko Takefusa, Hidemoto Nakada, Ryousei Takano, Tomohiro Kudoh, and Yoshio Tanaka</i>	
An Efficient Orchestration Engine for the Cloud	711
<i>Rafael Z. Frantz, Rafael Corchuelo, and José L. Arjona</i>	
Dynamic Data Redistribution for MapReduce Joins	717
<i>Steven Lynden, Yusuke Tanimura, Isao Kojima, and Akiyoshi Matono</i>	
Job Aware Scheduling Algorithm for MapReduce Framework	724
<i>Radheshyam Nanduri, Nitesh Maheshwari, A. Reddyraja, and Vasudeva Varma</i>	
Web Farm-inspired Computational Cluster in the Cloud	730
<i>Justin L. Rice, Vir V. Phoha, Pat Cappelaere, and Dan Mandl</i>	
Snow White Clouds and the Seven Dwarfs	738
<i>Stephen C. Phillips, Vegard Engen, and Juri Papay</i>	
Investigating the Impact of Deployment Configuration and User Demand on a Social Network Application in the Amazon EC2 Cloud	746
<i>Matheus Cunha, Nabor Mendonça, and Américo Sampaio</i>	
Market Implementation of Cloud Interoperability and Portability Research in IaaS and PaaS	
Cloud Computing Interoperability: The State of Play	752
<i>Nikolaos Loutas, Eleni Kamateri, Filippo Bosi, and Konstantinos Tarabanis</i>	
Cloud Application Monitoring: The mOSAIC Approach	758
<i>Massimiliano Rak, Salvatore Venticinque, Tamás Máhr, Gorka Echevarria, and Gorka Esnal</i>	
Emerging Issues on Cloud Computing	
End-to-End Policy-Based Encryption and Management of Data in the Cloud	764
<i>Siani Pearson, Marco Casassa Mont, Liqun Chen, and Archie Reed</i>	
Monitoring Intrusions and Security Breaches in Highly Distributed Cloud Environments	772
<i>Aryan Taheri Monfared and Martin Gilje Jaatum</i>	
Evaluating the Performance and Power Consumption of Systems with Virtual Machines	778
<i>Ricardo Lent</i>	
Information Stewardship in Cloud Ecosystems: Towards Models, Economics, and Delivery	784
<i>Adrian Baldwin, David Pym, Martin Sadler, and Simon Shiu</i>	

Security APIs for My Private Cloud - Granting Access to Anyone, from Anywhere at any Time	792
<i>David W. Chadwick and Matteo Casenove</i>	
Automatic Creation of an Ontological Knowledge Base from Grid and Cloud-based Wikipages	799
<i>Chaitali Gupta and Madhusudhan Govindaraju</i>	
Author Index	805