1. Record Nr. UNISA996464438703316

Titolo	Cloud computing and services science : 10th international conference, CLOSER 2020, Prague, Czech Republic, May 7-9, 2020, revised selected papers / / edited by Donald Ferguson, Claus Pahl, and Markus Helfert
Pubbl/distr/stampa	Cham, Switzerland : , : Springer, , [2021] ©2021
ISBN	3-030-72369-0
Descrizione fisica	1 online resource (357 pages) : illustrations
Collana	Communications in Computer and Information Science ; ; v.1399
Disciplina	004.6782
Soggetti	Cloud computing
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	Includes index.
Nota di contenuto	Intro -- Preface -- Organization -- Contents -- Trusted Client-Side Encryption for Cloud Storage -- 1 Introduction -- 1.1 Contribution -- 1.2 Outline -- 2 Background -- 2.1 Cloud Storage -- 2.2 File and Disk Encryption -- 2.3 Cryptomator -- 2.4 Intel Software Guard Extensions Technology -- 3 Related Work -- 3.1 Cloud Storage Security and Privacy -- 3.2 File and Disk Encryption with Intel SGX -- 3.3 Cloud Storage with Intel SGX -- 4 Proposal -- 5 Performance Evaluation -- 5.1 Experimental Setup -- 5.2 Methodology -- 5.3 File Writing Performance -- 5.4 File Reading Performance -- 6 Security Assessment -- 6.1 Threat Model -- 6.2 Intel SGX Security Assessment -- 6.3 Cryptomator-SGX Security Assessment -- 7 Limitations -- 8 Future Work -- 9 Conclusion -- References -- Relevance of Near-Term Quantum Computing in the Cloud: A Humanities Perspective -- 1 Introduction -- 2 Programming Model -- 2.1 Quantum Registers -- 2.2 Quantum Operations -- 2.3 Quantum Algorithms -- 2.4 Quantum Software Stack -- 2.5 Sample Research Questions -- 3 Quantum as a Service -- 3.1 Tooling -- 3.2 Deployment and Quantum Application as a Service -- 3.3 Sample Research Questions -- 4 Removing Hardware Dependencies -- 4.1 Problem -- 4.2 Hardware-Independent Processing -- 4.3 NISQ Analyzer -- 4.4 Quantum Compiler -- 4.5 Optimization of Quantum Algorithms -- 4.6 Sample Research Questions -- 5 Quantum Humanities: A Use Case

from Quantum Machine Learning -- 5.1 Quantum Humanities -- 5.2 MUSE -- 5.3 Potential Improvements -- 5.4 Categorical Data -- 5.5 Distance Measures and Feature Extraction -- 5.6 Sample Research Questions -- 6 Collaborative Quantum Application Platform -- 6.1 Key Entities -- 6.2 Requirements -- 6.3 Architecture -- 6.4 Sample Research Questions -- 7 Conclusions -- References.
A Framework for Comparative Evaluation of High-Performance Virtualized Networking Mechanisms -- 1 Introduction -- 1.1 Contributions -- 2 Background -- 2.1 Kernel-Based Networking and VNFs -- 2.2 Bypassing the Kernel's Networking Stack -- 2.3 Inter-container Communications with Kernel Bypass -- 2.4 High-Performance Switching Among Containers -- 3 Proposed Framework -- 4 Experimental Results -- 4.1 Testing Parameters -- 4.2 Kernel-Based Networking -- 4.3 Throughput Evaluations -- 4.4 Throughput Scalability Evaluations -- 4.5 Latency Performance Evaluations -- 4.6 Performance and Computational Requirements -- 5 Related Work -- 6 Conclusions and Future Work -- References -- Cloud Computing for Enabling Big Data Analysis -- 1 Introduction -- 2 Data Mining Cloud Framework (DMCF) -- 2.1 Workflow Formalisms -- 2.2 Workflow Examples -- 2.3 Workflow Study Cases -- 3 Parallel Social Data Analysis (ParSoDA) -- 3.1 Reference Architecture and Execution Flow -- 3.2 Usability and Scalability Evaluation -- 4 Nubytics -- 4.1 Architecture -- 4.2 Services -- 5 SMA4TD -- 5.1 Steps 1-2: Definition of Events and Places-of-Interest -- 5.2 Steps 3-4-5: Collection and Pre-processing of Geotagged Items, Identification of Users and Creation of the Input Dataset -- 5.3 Step 6-7: Data Mining and Results Visualization -- 5.4 Study Cases: FIFA World Cup 2014 and EXPO 2015 -- 6 RoI Mining -- 6.1 Methodology -- 7 Iterative Opinion Mining Using Neural Networks -- 7.1 Case Study: The 2016 US Presidential Election -- 8 Conclusions -- References -- A Multi-cloud Parallel Selection Approach for Unlinked Microservice Mapped to Budget's Quota: The PUM2Q -- 1 Introduction -- 2 Example for PUM2Q -- 3 Formulation -- 3.1 Availability Requirement -- 3.2 Response Time Requirement -- 3.3 Cost Requirement -- 4 PUM2Q Approach -- 4.1 First Level: Determine Candidate Services for a Microservice. 4.2 Second Level: Determine a Candidate Combination from All Providers -- 4.3 Third Level: Determine a Combination for a Microservice -- 5 Performance Assessment -- 5.1 Tool Description -- 5.2 Setting Scenarios -- 5.3 Experiments for PUM2Q -- 5.4 Experiments for UM2Q -- 5.5 Comparison Between PUM2Q and UM2Q -- 6 Images/Related Work -- 7 Conclusion -- References -- Live Migration Timing Optimization Integration with VMware Environments -- 1 Introduction -- 2 Background -- 2.1 Live Migration Configurations -- 2.2 VMware Virtual Networking Configuration -- 3 Cost Prediction Algorithm and Timing Optimization -- 3.1 The Research Problem -- 3.2 Cost Prediction Algorithm -- 3.3 Timing Optimization Algorithm -- 4 Integration with VMware vSphere -- 5 Testing Results -- 6 Conclusion -- References -- Using Self-Organizing Maps for the Behavioral Analysis of Virtualized Network Functions -- 1 Introduction -- 1.1 Contributions -- 1.2 Paper Organization -- 2 Related Work -- 3 Background Concepts -- 3.1 Self-Organizing Maps -- 3.2 VMware vRealize Operations Manager -- 4 Proposed Approach -- 4.1 Workflow -- 4.2 SOM Implementation -- 4.3 Hierarchical Grouping -- 4.4 Alerting -- 5 Experimental Results -- 5.1 Multi-metric Analysis -- 5.2 Hyper-parameters Grid Search -- 5.3 Per-VNF Analysis -- 5.4 Hierarchical Grouping -- 5.5 Alerting -- 6 Conclusions -- References -- From DevOps to NoOps: Is It Worth It? -- 1 Introduction -- 2 Background -- 2.1 DevOps-Based Cloud Model -- 2.2 NoOps-Based

Cloud Model -- 3 Methodology -- 3.1 Benchmarks -- 3.2 Deployment Strategies -- 3.3 Load Test Settings and Infrastructure -- 4 Results -- 4.1 Cinema Application -- 4.2 Primes-Python Function -- 4.3 Image-Processing Function -- 4.4 Nodeinfo Function -- 5 Discussion -- 6 Related Work -- 7 Conclusion -- 8 Availability -- References.
Fuzzy Container Orchestration for Self-adaptive Edge Architectures -- 1 Introduction -- 2 Background Concepts and Technologies -- 2.1 Platform and Architecture Concepts -- 2.2 Self-adaptiveness and Fuzziness Concepts -- 2.3 Platform Tools and Technologies -- 2.4 Towards a Lightweight Edge Platform -- 3 Platform and Controller Architecture -- 3.1 Platform and Cluster Configuration -- 3.2 Platform Monitoring -- 4 A Fuzzy Auto-scaling Controller -- 4.1 Principles of Auto-scaling -- 4.2 The Fuzzy Service -- 4.3 Controller Calibration and Configuration -- 4.4 Continuous Auto-scaling -- 5 Controller Evaluation -- 5.1 Evaluation Objectives and Metrics -- 5.2 Evaluation Setup -- 5.3 Evaluation Implementation -- 5.4 Discussion of Results -- 6 Related Work -- 7 Conclusions -- References -- Performance Management in Clustered Edge Architectures Using Particle Swarm Optimization -- 1 Introduction -- 2 Principles of Particle Swarm Optimization -- 2.1 Particle Swarm Optimization Basics -- 2.2 Basic Algorithm Definitions -- 2.3 Population Generation -- 2.4 Selection of Primary Nodes -- 2.5 Evaluation of the Objective Function -- 2.6 Recording the Best Position -- 3 Performance Optimization for Clustered Architectures -- 3.1 Edge Request Management -- 3.2 Orchestration Controller Design -- 3.3 Performance Optimization -- 3.4 PSO-Based Performance Optimization -- 3.5 Algorithm Definition (Pseudocode) -- 4 Framework Implementation and Evaluation -- 4.1 Test Environment and Objectives -- 4.2 Definition of PSO and BAT Parameters -- 4.3 Comparison of Fitness Values -- 4.4 Scenario-Based Comparison: Overview -- 4.5 Scenario 1: Request Variation -- 4.6 Scenario 2: Edge Node Variation -- 4.7 Scenario 3: Iteration Variation -- 5 Related Work -- 6 Conclusions and Future Work -- References -- Investigating IoT Application Behaviour in Simulated Fog Environments. 1 Introduction -- 2 Related Approaches in Fog-Cloud Simulation -- 3 Fog Modelling in Two Simulators -- 4 IoT Application Scenarios and Architecture Configuration for the Evaluation -- 5 Results of the First Evaluation Round -- 6 Further Analysis of the Investigated Simulators -- 7 The Second Evaluation Round -- 8 Result of the Third Evaluation Round -- 9 Conclusions -- References -- A Case for User-Defined Governance of Pure Edge Data-Driven Applications -- 1 Introduction -- 2 Community-Governed Services Model -- 2.1 Definition -- 2.2 System Model -- 2.3 Application Requirements -- 3 Case Study -- 4 Materials and Methods -- 4.1 Data Source -- 4.2 Simulation Model -- 4.3 Defining Trust Relationship -- 4.4 Experimental Design -- 4.5 Evaluation Metrics -- 5 Results and Discussion -- 5.1 Best Case Analysis: No Restrictions on Data Sharing -- 5.2 Analysis of Data Sharing Using Synthetic Social Network Restrictions -- 5.3 Analysis of Data Sharing Using Real Social Network Restrictions -- 6 Related Work -- 7 Conclusions -- References -- Cluster-Agnostic Orchestration of Containerised Applications -- 1 Introduction -- 2 Related Work -- 3 The TOSCA Specification -- 4 Analysis of Cluster Orchestrators -- 5 Achieving Cluster Interoperability -- 6 System Design and Implementation -- 6.1 Framework Architecture -- 6.2 YAML Parsing -- 6.3 BPMN Plans -- 6.4 Service Connectors -- 6.5 Prototype Implementation -- 7 Experiments -- 7.1 Use Case -- 7.2 Performance Evaluation -- 8 Conclusion -- References -- Self-healing in the Scope of Software-Based Computer and Mobile Networks -- 1 Introduction -- 2 Background -- 2.1 Fault Tolerance in the Data Plane -- 2.2 Fault

Tolerance in the Control and Management Planes -- 2.3 Related Work -- 3 Self-healing Framework -- 3.1 Framework Architecture -- 3.2 Monitoring of Different Network Layers.
3.3 Recovery by Detection and Prediction.