

CLOSER 2015 Abstracts

Area 1 - Cloud Computing Fundamentals

Full Papers Short Papers Posters

Area 2 - Services Science Foundation for Cloud Computing

Short Papers Posters

Area 3 - Cloud Computing Platforms and Applications

Full Papers Short Papers Posters

Area 4 - Cloud Computing Enabling Technology

Full Papers Short Papers Posters

Area 5 - Mobile Cloud Computing and Services

Full Papers Short Papers

Area 1 - Cloud Computing Fundamentals

Full Papers

Paper Nr: 12

Title: Effects of Active Cooling on Workload Management in High Performance Processors

Authors: Won Ho Park and C. K. Ken Yang

Abstract: This paper presents an energy-efficient workload scheduling methodology for multi-core multi-processor systems under actively cooled environment that improves overall system power performance with minimal response time degradation. Using a highly efficient miniature-scale refrigeration system, we show that active-cooling by refrigeration on a per-server basis not only leads to substantial power-performance improvement, but also improves the overall system performance without increasing the overall system power including the cost of cooling. Based on the measured results, we present a model that captures different relations and parameters of multi-core processor and the refrigeration system. This model is extended to illustrate the potential of power optimization of multi-core multi-processor systems and to investigate different methodologies of workload scheduling under the actively cooled environment to maximize power efficiency while minimizing response time. We propose an energy-efficient workload scheduling methodology that results in total consumption comparable to the spatial subsetting scheme but with faster response time under the actively cooled environment. The actively cooled system results in $\geq 29\%$ of power reduction over the non-refrigerated design across the entire range of utilization levels. The proposed methodology is further combined with the G/G/m-model to investigate the trade-off between the total power and target SLA requirements.

Paper Nr: 24

Title: A Mathematical Programming Approach to Multi-cloud Storage

Authors: Makhlof Hadji

Abstract: This paper addresses encrypted data storage in multi-cloud environments. New mathematical models and algorithms are introduced to place and replicate encrypted data chunks and ensure high availability of the data. To enhance data availability, we present two cost-efficient algorithms based on a complete description of a linear programming approach of the multi-cloud storage problem. Performance assessment results, using simulations, show the scalability and cost-efficiency of the proposed multi-cloud distributed storage solutions.

Paper Nr: 47

Title: Cloud Provider Transparency - A View from Cloud Customers

Authors: Daniela Cruzes and Martin Gilje Jaatun

Abstract: A major feature of public cloud services is that data are processed remotely in unknown systems that the users do not own or operate. This context creates a number of challenges related to data privacy and security and may hinder the adoption of cloud technology. One of these challenges is how to maintain transparency of the processes and procedures while at the same time providing services that are secure and cost effective. This paper presents results from an empirical study in which the cloud customers identified a number of transparency requirements to the adoption of cloud providers. We have compared our results with previous studies, and have found that in general, customers are in synchrony with research criteria for cloud service provider transparency, but there are also some extra pieces of information that customers are looking for.

Paper Nr: 81

Title: OCCl and TTCN-3 - Towards a Standardized Cloud Quality Assessment Framework

Authors: Yongzheng Liang

Abstract: Impacting basically all types of IT infrastructures The Cloud is one of the most important evolving IT paradigms. A standard-based Cloud quality and compliance assessment framework will be therefore of utmost importance. Bringing together the Open Cloud Computing Interface OCCl and the ETSI standardized test specification language TTCN-3 and related test methodologies this paper is going to demonstrate initial steps towards such a framework. Taking into account the diversity of Cloud infrastructures, of service providers, and related architectural, harmonization and standardization effort this approach is mainly motivated by studying Cloud-related effort of the NIST Cloud Computing Program and the ETSI Cloud Standards Coordination (CSC). Reflecting the "Cloudiness" of the Software Defined Network (SDN) and ETSI Network Functions Virtualization (NFV) this paper is considering these initiatives as necessary elements of the scope of every future standardized Cloud quality assessment framework as well.

Paper Nr: 84

Title: Using Cloud-Aware Provenance to Reproduce Scientific Workflow Execution on Cloud

Authors: Khawar Hasham, Kamran Munir and Richard McClatchey

Provenance has been thought of a mechanism to verify a workflow and to provide workflow reproducibility. This provenance of scientific workflows has been effectively carried out in Grid based scientific workflow systems. However, recent adoption of Cloud-based scientific workflows present an opportunity to investigate the suitability of existing approaches or propose new approaches to collect provenance information from the Cloud and to utilize it for workflow reproducibility on the Cloud infrastructure. This paper presents a novel approach that can assist in mitigating this challenge. This approach can collect Cloud infrastructure information along

Abstract: with workflow provenance and can establish a mapping between them to provide a Cloud-aware provenance. The reproducibility of the workflow execution is performed by: (a) capturing the Cloud infrastructure information (virtual machine configuration) along with the workflow provenance, (b) re-provisioning the similar resources on the Cloud and re-executing the workflow on them and (c) by comparing the outputs of workflows. The evaluation of the prototype suggests that the proposed approach is feasible and can be investigated further. Since there is no reference model for workflow reproducibility on Cloud exists in the literature, this paper also attempts to present a model that is used in the proposed design to achieve workflow reproducibility in the Cloud environment.

Paper Nr: 104

Title: Addressing Issues of Cloud Resilience, Security and Performance through Simple Detection of Co-locating Sibling Virtual Machine Instances

Authors: John O'Loughlin and Lee Gillam

Most current Infrastructure Clouds are built on shared tenancy architectures, with resources shared amongst large numbers of customers. However, multi tenancy can lead to performance issues (so-called "noisy neighbours") and also brings potential for serious security breaches such as hypervisor breakouts. Consequently, there has been a focus in the literature on identifying co-locating instances that are being affected by noisy neighbours or suggesting that such instances are vulnerable to attack. However, there is

Abstract: limited evidence of any such attacks in the wild. More beneficially, knowing that there is co-location amongst your own Virtual Machine instances (siblings) can help to avoid being your own worst enemy: avoiding your instances acting as your own noisy neighbours, building resilience through ensuring host-based redundancy, and/or reducing exposure to a single compromised host. In this paper, we propose and demonstrate a simple test to detect co-locating sibling instances on Xen-based Clouds, as could help address such needs, and evaluate its efficacy on Amazon's EC2

Short Papers

Paper Nr: 16

Title: P-TOSCA Portability of SOA Applications

Authors: Marjan Gusev, Magdalena Kostoska, Sasko Ristov and Aleksandar Donevski

Even more frequently, the customers express their increasing need to change the cloud provider and/or the operating cloud environment in order to avoid vendor lock-in. We analyze portability as the transferability of an application from on-premise onto a

Abstract: cloud (migration) and among different clouds (porting). The contribution of this paper is twofold: 1) demonstration of the P-TOSCA model for automated migration and porting of SOA applications onto a cloud and/or switch between cloud providers, and 2) evaluation of a significant time reduction in migration and porting.

Paper Nr: 21

Title: A Cloud-based Data Analysis Framework for Object Recognition

Authors: Rezvan Pakdel and John Herbert

This paper presents a Cloud-based framework using parallel data processing to identify and recognize an object from an image. Images contain a massive amount of information. Features such as shape, corner, color, and edge can be extracted from images.

Abstract: These features can be used to recognize an object. In a Cloud-based data analytics framework, feature detection algorithms can be done in parallel to get the result faster in comparison to a single machine. This study provides a Cloud-based architecture as a solution for large-scale datasets to decrease processing time and save hardware costs. The evaluation results indicate that the proposed approach can robustly identify and recognize objects in images.

Paper Nr: 25

Title: Factors Affecting Cloud Adoption and Their Interrelations

Authors: Radhika Garg and Burkhard Stiller

Cloud Computing has emerged as a paradigm that relies on sharing resources over the network and, therefore, potentially has cost advantages in terms of lower variable and capital cost. However, the adoption of cloud-based technology for a given IT (Information Technology) setting is a complex decision as it is influenced by multiple interdependent factors. To successfully adopt cloud-based

Abstract: services and evaluate their consequential impact, relevant factors, which denote the performance of such services, have to be identified. This paper, therefore, analyzes and identifies relevant technical, economical, and organizational factors. This is performed as exploratory research consisting of performing (a) a literature review and (b) multiple case-studies with 17 organizations, who have adopted or plan to adopt cloud-based services. Also, as these factors are not mutually exclusive, this paper discusses interrelations of these factors and its complexity.

Paper Nr: 48

Title: Quality of Service Trade-offs between Central Data Centers and Nano Data Centers

Authors: Farzaneh Akhbar and Tolga Ovatman

Nano data centers are one of the latest trends in cloud computing aiming towards distributing the computing power of massive data centers among the clients in order to overcome setup and maintenance costs. The distribution process is done over the already present computing elements in client houses such as tv receivers, wireless modems, etc. In this paper we investigate the feasibility of using nano data centers instead of conventional data centers containing accumulated computing power. We try to draw the lines that may affect the decision of nano data center approach considering important parameters in cloud computing such as memory capacity,

Abstract: diversity of user traffic and computing costs. We also investigate the thresholds for these parameters to find out the conditions that make more sense to set up nano data centers as the best replacement of Central Data Centers. We use a CloudSim based simulator, namely CloudAnalyst, for Data Center performance experiments in java. Our results show that 1 gigabyte memory capacity can be seen as a threshold for response time improvement of nano data centers. For nano data centers with more memory capacity there will not be any improvement in response times that leverages the performance cost. We also combine the results of response time and performance cost to provide a similar threshold.

Paper Nr: 56

Title: Cloud Readiness Assessment of Legacy Application

Authors: Flavio Corradini, Francesco De Angelis, Andrea Polini and Samuele Sabbatini

Applications and services hosted in the cloud are increasing continuously. Cloud technology offers important perspectives (performance, high availability, elasticity) and it enables new business models. Unfortunately, this new paradigm faces unprecedented requirements not addressed in legacy application (multi-tenancy, scalability, etc.). This leads to complex re-engineering phases in

Abstract: order to migrate existing software into a cloud environment. Before starting a migration, it is important to analyze the cloud compliance of the application, what to expect after the migration and the effort required to fulfill these expectations. This paper assesses a way to extract an index that describes the feasibility of the re-engineering. We test the metric with a real application that needs to be migrated to a private cloud.

Paper Nr: 60

Title: Development of an Anything Relationship Management Prototype for the Smart Factory

Authors: Jonathan Knoblauch, Rebecca Bulander and Thomas Greiner

The Internet of Things, Services and Data (IoTSD) enters into more and more areas of the business, private and public sector. Typical areas are Smart Factory, Smart Home, Smart Grid, Connected Vehicles and Smart City. The area of Smart Factory (also called industrial internet) will be the most important one in the manufacturing sector. For several years there has been another development in information and communications technology (ICT) observable, called Anything Relationship Management (xRM), trying to systematically manage all stakeholders, physical objects and virtual entities of an enterprise through the use of powerful IT platforms.

Abstract: xRM can be used as a cloud management platform for smart industrial production units combined with stakeholder management. In this paper we use xRM to develop a top-down prototype in the Smart Factory environment. The main objective is to demonstrate how xRM could be used in the future Smart Factory. We therefore recreate the structure of an existing machine for mixing liquids as a service on an xRM cloud platform. Furthermore typical data exchange activities between machine and an xRM cloud platform as well as customers and production machine are simulated. The xRM prototype demonstrates why using an xRM cloud platform is helpful for flexible production environments.

Paper Nr: 66

Title: Redefining the Cloud based on Beneficial Service Characteristics - A New Cloud Taxonomy Leads to Economically Reasonable Semi-cloudification

Authors: Bastian Kemmler and Dieter Kranzlmüller

Cloud services promise benefits for customers and providers such as scalability, elasticity and reduced investment costs. Unfortunately, many of the promised benefits are not fulfilled by today's cloud offerings and not every service can be cloudified, e.g. if the service's intrinsic structure contains unavoidable time-consuming or manual tasks. A new cloud definition, based on a survey and comparison of existing cloud definitions, but derived from beneficial cloud characteristics, leads to a service-oriented understanding of

Abstract: clouds and provides an extension to the usual cloud service types. The characteristics of the given cloud definition uncover the so-called "MOUSETRAPS" of cloud services. The term "semi-cloudification" for the transformation of services towards a cloudified state presents a solid foundation for further discussions on the topic and enables the improvement of non-cloudifiable services by semi-cloudification. Even services which partly consist of unavoidable time-consuming or manual tasks qualify for semi-cloudification.

Paper Nr: 69

Title: CoMA: Resource Monitoring of Docker Containers

Authors: Lara Lorna Jiménez, Miguel Gómez Simón, Olov Schelén, Johan Kristiansson, Kåre Synnes and Christer Åhlund

This research paper presents CoMA, a Container Monitoring Agent, that oversees resource consumption of operating system level virtualization platforms, primarily targeting container-based platforms such as Docker. The core contribution is CoMA, together with a quantitative evaluation verifying the validity of the measurements reported by the agent for three metrics: CPU, memory and block I/O.

Abstract: The proof-of-concept is implemented for Docker-based systems and consists of CoMA, the Ganglia Monitoring System and the Host sFlow agent. This research is in line with the rising trend of container adoption which is due to the resource efficiency and ease of deployment. These characteristics have set containers in a position to topple virtual machines as the reigning virtualization technology in data centers.

Paper Nr: 89

Title: Towards Dynamic QoS Monitoring in Service Oriented Architectures

Authors: Norman Ahmed and Bharat Bhargava

Service Oriented Architecture (SOA) is an architectural style that provides agility to align technical solutions to modular business Web Services (WS) that are well decoupled from their consumers. This agility is established by interconnecting WS family of standards specification protocols (commonly referred to as WS-* (WS-star)) to enable security, ease of service interoperability and orchestration complexities when extending services across organizational boundaries. While orchestrating services or chaining services in varying ways to satisfy different business needs, on highly scalable cloud platforms is undeniably useful, it is increasingly challenging to

Abstract: effectively monitor Quality of Service (QoS), especially, service response time. This is due to a) lack of proper formulation of the WS-star interconnections mechanisms, and b) the transient performance behaviour intrinsic to the heterogeneity of the hardware and shared virtualized network and IO resources built on the cloud platforms. We present an analysis of WS-star standards, classifying and discussing their inter-dependencies to provide a basis for QoS monitoring context on protocol formulation. We then illustrate a practical implementation of a dynamic QoS monitoring mechanism using runtime service instrumentation with Aspect Oriented Programming (AOP). Preliminary evaluations show the efficiency of computing QoS on a transient performance cloud platform.

Paper Nr: 102

Title: Offline Scheduling of Map and Reduce Tasks on Hadoop Systems

Authors: Aymen Jlassi, Patrick Martineau and Vincent Tkindt

MapReduce is a model to manage quantities massive of data. It is based on the distributed and parallel execution of tasks over the cluster of machines. Hadoop is an implementation of MapReduce model, it is used to offer BigData services on the cloud. In this paper, we expose the scheduling problem on Hadoop systems. We focus on the offline-scheduling, expose the problem in a mathematic model and use the timeindexed formulation. We aim consider the maximum of constraints of the MapReduce environment. Solutions for the presented model would be a reference for the on-line Schedules in the case of low and medium instances. Our work is useful in term of the problem definition: constraints are based on observations and take into account resources consumption, data locality, heterogeneous machines and workflow management; this paper defines boundaries references to evaluate the online model.

Abstract:

Paper Nr: 105

Title: A Generalized Service Replication Process in Distributed Environments

Authors: Hany F. El Yamany, Marwa F. Mohamed, Katarina Grolinger and Miriam A. M. Capretz

Replication is one of the main techniques aiming to improve Web services' (WS) quality of service (QoS) in distributed environments, including clouds and mobile devices. Service replication is a way of improving WS performance and availability by creating several copies or replicas of Web services which work in parallel or sequentially under defined circumstances. In this paper, a generalized replication process for distributed environments is proposed based on established replication studies. The proposed replication

Abstract: process consists of three main steps: sensing the environment characteristics, determining the replication strategy, and implementing the selected replication strategy. To demonstrate application of the proposed generalized replication process, a case study in the telecommunication domain is presented. The adequacy of the selected replication strategy is demonstrated by comparing it to another replication strategy as well as to a non-replicated service. The authors believe that the proposed replication process will help service providers to enhance QoS and accordingly attract more customers.

Paper Nr: 114

Title: User Requirement and Behavioral Aspects in Web Service Discovery

Authors: Wala Ben messaoud, Khaled Ghédira and Youssef Ben Halima

In web service (WS) discovery, behavioral aspect has been defined as the sequence of WS operations. The motivation to introduce the behavioral aspect is to offer to the consumer the possibility to choose his WS according to his requirements. The aim is to include the execution manner of WS operations as a new criterion and to apply a selection method if more than one WS candidate is filtered.

Abstract: In this paper, we envision to implement WS discovery approach based on behavioral aspects to fulfill the selection of the precise execution order. This approach ensures an execution order of operations in accordance to consumer needs. The execution manner criterion is defended by state chart as a scheduling method and WordNnet as a lexical database. Moreover, semantic equivalences have to be considered in order to solve equivalence between many WS candidates which satisfy consumer needs.

Paper Nr: 115

Title: PaaSword: A Holistic Data Privacy and Security by Design Framework for Cloud Services

Authors: Yiannis Verginadis, Antonis Michalas, Panagiotis Gouvas, Gunther Schiefer, Gerald Hübsch and Iraklis Paraskakis

The valuable transformation of organizations that adopt cloud computing is indisputably accompanied by a number of security threats that should be considered. In this paper, we outline significant security challenges presented when migrating to a cloud environment and propose PaaSword – a novel holistic, data privacy and security by design, framework that aspires to alleviate them. The

Abstract: envisaged framework intends to maximize and fortify the trust of individual, professional and corporate users to cloud services. Specifically, PaaSword involves a context-aware security model, the necessary policies enforcement and governance mechanisms along with a physical distribution, encryption and query middleware, aimed at facilitating the implementation of secure and transparent cloud-based applications.

Paper Nr: 116

Title: Classifying Security Threats in Cloud Networking

Authors: Bruno M. Barros, Leonardo H. Iwaya, Marcos A. Simplicio Jr., Tereza C. M. B. Carvalho, András Méhes and Mats Näslund

A central component of managing risks in cloud computing is to understand the nature of security threats. The relevance of security concerns are evidenced by the efforts from both the academic community and technological organizations such as NIST, ENISA and CSA, to investigate security threats and vulnerabilities related to cloud systems. Provisioning secure virtual networks (SVNs) in a multi-tenant environment is a fundamental aspect to ensure trust in public cloud systems and to encourage their adoption. However,

Abstract: comparing existing SVN-oriented solutions is a difficult task due to the lack of studies summarizing the main concerns of network virtualization and providing a comprehensive list of threats those solutions should cover. To address this issue, this paper presents a threat classification for cloud networking, describing threat categories and attack scenarios that should be taken into account when designing, comparing, or categorizing solutions. The classification is based on the CSA threat report, building upon studies and surveys from the specialized literature to extend the CSA list of threats and to allow a more detailed analysis of cloud network virtualization issues.

Paper Nr: 123

Title: Setting Priorities - A Heuristic Approach for Cloud Data Center Selection

Authors: Ronny Hans, David Steffen, Ulrich Lampe, Björn Richerzhagen and Ralf Steinmetz

A rising number of multimedia applications with Quality of Service requirements is delivered via cloud computing platforms. To reduce latencies between data centers and customers, providers need to enhance and utilize their cloud infrastructure by providing resources closer to the consumer. For planning such infrastructures and efficiently assigning existing resources, capable algorithms to solve the

Abstract: underlying optimization problem are required. With our priority-based heuristic approach, we are able to reduce the computation time by up to 99.99% compared to an exact approach, while retaining a favorable solution quality.

Posters

Paper Nr: 29

Title: A Comparative Study of Current Open-source Infrastructure as a Service Frameworks

Authors: Theo Lynn, Graham Hunt, David Corcoran, John Morrison and Philip Healy

With the growth of cloud computing in recent years, several commercial and open source IaaS frameworks have emerged. The development of open source IaaS solutions offers a free and flexible alternative to commercial cloud services. The main contribution of this paper is to provide a qualitative comparative of current open-source IaaS frameworks. Existing research papers examining open

Abstract: source IaaS frameworks have focused on comparing OpenStack with a small number of alternatives. However, current research fails to adequately compare all major open source frameworks in a single study and notably lacks the inclusion of CloudStack. Our research paper provides the first overview of the five main open source cloud IaaS frameworks – OpenStack, CloudStack, OpenNebula, Eucalyptus and Nimbus. As such, this review provides researchers and potential users with an up to date and comprehensive overview of the features of each solution and allows for an easy comparison between the open source solutions.

Paper Nr: 45

Title: CSP Formulation for Scheduling Independent Jobs in Cloud Computing

Authors: M'hamed Mataoui, Faouzi Sebbak, Kadda Beghdad Bey and Farid Benhammedi

This paper investigates the use of Constraint Satisfaction Problem formulation to schedule dependent jobs in heterogeneous cloud environment. Our formulation provides a basis for computing an optimal Makespan using job and machine reordering heuristics based on Min-min algorithm result. The combination of these heuristics with the weighted constraints allows improving the efficiency of the tree search algorithm to schedule jobs with considerable space search reduction. The proposed CSP model is validated through simulation experiments against clusters of 10 virtual machines. The results demonstrate that our model is able to efficiently allocate resources for jobs with significant performance gains between 18% - 40% compared to the Min-Min heuristic results to optimize the Makespan.

Paper Nr: 70

Title: A Survey of Trust Management Models for Cloud Computing

Authors: Flavio Corradini, Francesco De Angelis, Fabrizio Ippoliti and Fausto Marcantoni

Abstract: Over the past few years, cloud computing has been widely adopted as a paradigm for large-scale infrastructures. In such a scenario, new security risks arise when different entities or domains share the same group of resources. Involved organizations need to establish some kind of trust relationships, able to define appropriate rules that can control which and how resources and services are going to be shared. The management of trust relationships represents a key challenge in order to meet high security requirements in cloud computing environments. This allows also to boost consumers confidence in cloud services, promoting its adoption. Establishing trust with cloud service providers supports to have confidence, control, reliability, and to avoid commercial issues like lock in. This paper proposes a survey of existing trust management models addressing collaboration agreements in cloud computing scenarios. Main limitations of current approaches are outlined and possible improvements are traced, as well as a future research path.

Paper Nr: 94

Title: Reality Vs Hype - Does Cloud Computing Meet the Expectations of SMEs?

Authors: Katie Wood and Kevan Buckley

Abstract: Small and Medium Enterprises (SMEs) have become a primary target audience for Cloud Service Providers (CSP), such as Amazon and Microsoft to promote their cloud offering. CSP strong marketing campaigns of 'promised' benefits from using their clouds is an attractive offer for SMEs especially where resources are limited and they wish to become more agile and reduce IT costing to be competitive with larger rivals. This paper argues that once SMEs remove the hype surrounding the concept of cloud computing (CC), the reality of significant benefits do not materialize for SMEs. This paper demonstrates, through working with SMEs considering the options of CC that the challenges and risks associated with cloud might actually hinder the business, rather than providing any real long term value.

Paper Nr: 111

Title: Implementation of Cloud ERP - Moderating Effect of Compliance on the Organizational Factors

Authors: Shivam Gupta and Subhas C. Misra

Abstract: Cloud ERP has changed the way business can be done for Small and Medium Enterprises (SMEs). The two important benefits offered by Cloud ERP are: (a) SMEs can log into the internet from any place to access applications and data services at any point in the time. (b) Pay for the services that are used or needed. Although Cloud ERP has taken the IT world by storm and with all the advancement that has taken place so far, there are still issues and challenges that require to be addressed. This paper relates issues pertaining to Compliance with Organizational factors for successful implementation of Cloud ERP.

Area 2 - Services Science Foundation for Cloud Computing

Short Papers

Paper Nr: 33

Title: Automated Mapping of Business Process Execution Language to Diagnostics Models

Authors: Hamza Ghandorh and Hanan Lutfiyya

Abstract: This paper illustrates how a specification of a business process can be automatically mapped to a fault diagnostic model. Observed failures at run-time are quickly analyzed through the diagnostic model to determine the faulty service.

Paper Nr: 54

Title: Context-aware Security@run.time Deployment

Authors: Wendpanga Francis Ouedraogo, Frederique Biennier, Catarina Ferreira Da Silva and Parisa Ghodous

Abstract: Taking advantage of the agility and interoperability provided by Service Oriented Architecture (SOA), Web 2.0 and XaaS (Anything as a Service) technologies, more and more collaborative Business Processes (BP) are set "on demand" by selecting, composing and orchestrating different business services depending on the current need. This involves re-thinking the way information, services and applications are organized, deployed, shared and secured among multi-cloud environment. Fitting this de-perimeterized and evolving execution context requires organising the service protection in a dynamic way in order to provide an up to date and consistent protection. To fit this goal, we propose to integrate the different protection requirements defined according to the business environment in a single security policy. Then we plug a context-aware security deployment architecture on the cloud service middleware to analyse both the security policy and the execution context to select, compose and orchestrate the convenient protection means. A proof of concept built on Frascati middleware is used to evaluate the impact of this "on-line" security mediation.

Paper Nr: 58

Title: Choreography-based Consolidation of Interacting Processes Having Activity-based Loops

Authors: Sebastian Wagner, Oliver Kopp and Frank Leymann

Abstract: Choreographies describe the interaction between two or more parties. The interaction behavior description might contain loops. In case two parties want to merge their behavior to gain competitive advantage, the contained loop constructs also have to be merged. This paper presents a language-independent discussion on loop-structure pairing in choreographies and possible merging strategies. Thereby, the focus is turned on loops grouping child activities to be iterated. We show the feasibility of the merging strategies by applying them to BPEL-based choreographies.

Paper Nr: 118

Title: BPMN Extensions for Decentralized Execution and Monitoring of Business Processes

Authors: Jonas Seeuww, Gregory Van Seghbroeck, Bruno Volckaert and Filip De Turck

Software-as-a-service (SaaS) providers are further expanding their offering by growing into the space of business process outsourcing (BPO). Therefore, the SaaS provider wants to administer and manage the business process steps according to a service level agreement. Outsourcing of business processes results in decentralized business workflows. However, current business process modeling languages, e.g. Business Process Execution Language (BPEL), Business Process Model and Notation (BPMN), are based highly on a centralized execution model and current BPMN engines offer limited constructs for federation and decentralized execution.

Abstract: To guarantee execution of business processes according to a service level agreement, different parties involved in a federated workflow must be able to inspect the state of external workflows. This requires advanced inspection interfaces and monitoring facilities. Current business process modeling languages must thus be extended to support monitoring in the specification, support modeling and support deployment of decentralized workflows. In this paper, correlation and monitoring extensions for BPMN are described. These extensions to BPMN are described such that the existing specification can still be used as is in a backwards compatible way.

Paper Nr: 129

Title: A Smart Decisional Cognitive System based on Self-adaptability of Web Services to the Context

Authors: Faiçal Felhi, Marwa Ayadi and Jalel Akaichi

Abstract: Memory loss or cognitive stimulation application for handicapped people is the subject of a recent field of studies in a information systems. In this way, Web services are a solution for the integration of distributed information systems, autonomous, heterogeneous and auto adaptable to the context. In this paper, we are interested in defining a new solution for a smart and decisional cognitive system based on self-adaptability of Web services to the context and showing this solution by a case study.

Posters

Paper Nr: 19

Title: Business Process Generation by Leveraging Complete Search over a Space of Activities and Process Goals

Authors: Dipankar Deb, Nabendu Chaki and Aditya Ghose

Abstract: An efficient and flexible business process not only helps an organization to meet the requirements of the evolving surroundings but also may facilitate a competitive advantage over other companies towards delivering the desired services. This is even more critical for an emerging paradigm like cloud based deployment. In this paper, we introduce a novel mechanism to generate the business process suitable for specific organizations. The approach provides an automated way to build the possible business processes for a given set of tasks that fulfills the goal and satisfies the constraints of an organization. In step 1, we show how to generate the finite space of all possible designs for a given set of tasks. Secondly, we accumulate the effect of each step to deduce the final effect of each possible process design and to ensure that the redesigned set of steps still realizes the service goal. The designs not meeting the service goals are eliminated from the space. In step 3, the rest of the designs are checked for the constraint satisfaction subject to some specific cases. The framework provides a comprehensive, both syntactically and semantically correct, consistent business process generation methodology that adheres to the target business goals and constraints.

Paper Nr: 30

Title: "BPELanon" - Protect Business Processes on the Cloud

Authors: Marigianna Skouradaki, Vincenzo Ferme, Frank Leymann, Cesare Pautasso and Dieter H. Roller

Abstract: The advent of Cloud computing supports the offering of many Business Process Management applications on a distributed, per-use basis environment through its infrastructure. Due to the fact that privacy is still an open issue in the Cloud, many companies are reluctant to move their Business Processes on a public Cloud. Since the Cloud environment can be beneficiary for the Business Processes, the investigation of privacy issues needs to be further examined. In order to enforce the Business Process sharing on the Cloud we propose a methodology ("BPELanon") for the anonymization of Business Processes expressed in the Web Service Business Process Execution Language (BPEL). The method transforms a process to preserve its original structure and run-time behavior, while completely anonymizing its business semantics. In this work we set the theoretical framework of the method and define a five management layers architecture to support its realization. We developed a tool that implements the "BPELanon" method, validate its functionality and evaluate its performance against a collection of real-world process models that were conducted in the scope of research projects.

Paper Nr: 40

Title: Cross-layer Service Adaptation - State-of-the-Art, Shortcoming Analysis, and Future Research Directions

Authors: Ameni Meskini, Yehia Taher, Rafiqul Haque and Yahya Slimani

Abstract: In the past few years several cross-layer monitoring and adaptation technologies have been proposed. Although these are cross-layer adaptation technologies, however, in practice they focus on a particular layer. Some solutions involves two layers, yet none of the existing solutions do not consider all the layers during adaptation process. Furthermore, cross-layer adaptation approaches generate incompatibility problems. This is an adaptation coordination problem. Incompatibility refers to the situations where the adaptation is performed in a layer is not compatible with the constraints exposed by the other layers. This survey aims at studying and analyzing current approaches for web services adaptation, discussing their shortcomings and proposing research directions on cross-layer web service adaptation.

Paper Nr: 46

Title: The Influence of the Provider's Service Fairness on the Customer's Service Recovery Satisfaction and on Positive Behavioral Intentions in Cloud Computing

Authors: Montri Lawkobkit and Roland Blomer

Abstract: The study shows a statistically significant positive effect between the provider's perceived structural service fairness and the customer's service recovery satisfaction and, in turn, also shows statistically positive regression weights between the customer's service recovery satisfaction and the intension to react positively in three directions: (1) to continue with the software, (2) to propagate a positive word-of-mouth (WOM), (3) to give honest feedback. The influence of the provider's perceived social service fairness on the customer's service recovery satisfaction does not appear to be significant but indicates a positive correlation. The study is based on data collected via a structured questionnaire from qualified users who have subscribed to Business-to-Business customer relationship management software and who use it as Software-as-a-Service in the cloud. Structural Equation Modelling was applied for the data analysis in order to confirm the chosen dependency model. The findings may help service providers to better understand their customers and to stimulate constructive actions to their continual improvement process.

Paper Nr: 117

Title: Key Requirements for Predictive Analytical IT Service Management - Architectural Key Characteristics for a Cloud based Realization

Authors: Christopher Schwarz, Hans-Peter Bauer, Lukas Blödmann and Erwin Zinsler

While trying to maintain sustainable competitive advantage, IT service providers are challenged with tremendous service complexity and a low level of flexibility caused by the lack of transparency, constrained scalability and the missing ability to identify needed service measures proactively. For overcoming these challenges, this paper presents a well-evaluated set of identified key requirements for a feasible realization of a highly scalable cloud based architecture that supports predictive analytics in several domains of IT Service Management. This presented concept goes far beyond traditional approaches and pertinent state-of-the-art software solutions by focusing on business analyses based on knowledge creation and domain-independent knowledge sharing. The proposed approach is based on profound analyses of related work as well as modern service oriented design and business analyses paradigms. It provides semantic complexity handling, structured and multi-layered service interaction, cloud-enabled scalability management as well as predictive business analyses based on semantic reasoning, decision-making support and pattern recognition. The derived results eventually provide solution architects with a feasible and technical independent fundament for architectural implementation decisions. It ultimately enables IT service providers to cope with modern flexibility needs and complexity challenges and therefore to continuously satisfy customers to gain competitive advantage.

Abstract:

Area 3 - Cloud Computing Platforms and Applications

Full Papers

Paper Nr: 18

Title: Secure Evidence Collection and Storage for Cloud Accountability Audits

Authors: Thomas Ruebsamen, Tobias Pulls and Christoph Reich

Cloud accountability audits can be used to strengthen trust of cloud service customers in cloud computing by providing reassurance regarding the correct processing of personal or confidential data in the cloud. However, such audits require various information to be collected. The types of information range from authentication and data access logging to location information, information on security controls and incident detection. Correct data processing has to be proven, which immediately shows the need for secure evidence record storage that also scales with the huge number of data sources as well as cloud customers. In this paper, we introduce Insyndās a suitable cryptographic mechanism for storing evidence for accountability audits in our previously proposed cloud accountability audits architecture. We present our reasoning for choosing Insynd by showing a comparison of Insynd properties with requirements imposed by accountability evidence collection as well as an analysis how security threats are being mitigated by Insynd. Additionally, we describe an agent-based evidence collection process with a special focus on security and privacy protection.

Paper Nr: 52

Title: Supporting Multiple Persistence Models for PaaS Applications using MDE - Issues on Cloud Portability

Authors: Elias Adriano Nogueira da Silva, Daniel Lucrédio, Ana Moreira and Renata Fortes

In cloud computing, lock-in refers to the difficulty of porting an application from one platform to another. An example of such difficulty can be witnessed when porting an application from Platform-as-a-Service Google App Engine to Microsoft Azure. Differences in their implementations are substantial, yielding non-portable applications. Standardization could address this problem, but existing initiatives are still to be accepted. This paper addresses lock-in by proposing a model-driven engineering design approach that decouples platform specific code from the application logic. The resulting platform-independent models, as well as corresponding model transformations, can be reused to generate distinct platform-specific implementations, hence reducing the programming effort spent coding repetitive tasks. Such transformations can be made available for reuse on a repository for cloud providers. We have implemented transformations to handle persistence for Google App Engine and Azure, and discuss how model-driven engineering can reconcile the differences between features of the persistence models of GAE and Azure.

Paper Nr: 85

Title: A Lightweight Tool for Anomaly Detection in Cloud Data Centres

Authors: Sakil Barbhuiya, Zafeirios Papazachos, Peter Kilpatrick and Dimitrios S. Nikolopoulos

Cloud data centres are critical business infrastructures and the fastest growing service providers. Detecting anomalies in Cloud data centre operation is vital. Given the vast complexity of the data centre system software stack, applications and workloads, anomaly detection is a challenging endeavour. Current tools for detecting anomalies often use machine learning techniques, application instance behaviours or system metrics distribution, which are complex to implement in Cloud computing environments as they require training, access to application-level data and complex processing. This paper presents LADT, a lightweight anomaly detection tool for Cloud data centres that uses rigorous correlation of system metrics, implemented by an efficient correlation algorithm without need for training or complex infrastructure set up. LADT is based on the hypothesis that, in an anomaly-free system, metrics from data centre host nodes and virtual machines (VMs) are strongly correlated. An anomaly is detected whenever correlation drops below a threshold value. We demonstrate and evaluate LADT using a Cloud environment, where it shows that the hosting node I/O operations per second (IOPS) are strongly correlated with the aggregated virtual machine IOPS, but this correlation vanishes when an application stresses the disk, indicating a node-level anomaly.

Paper Nr: 88

Title: Performance and Cost Evaluation for the Migration of a Scientific Workflow Infrastructure to the Cloud

Authors: Santiago Gómez Sáez, Vasilios Andrikopoulos, Michael Hahn, Dimka Karastoyanova, Frank Leymann, Marigianna Skouradaki and Karolina Vukojevic-Haupt

The success of the Cloud computing paradigm, together with the increase of Cloud providers and optimized Infrastructure-as-a-Service (IaaS) offerings have contributed to a raise in the number of research and industry communities that are strong supporters of migrating and running their applications in the Cloud. Focusing on eScience simulation-based applications, scientific workflows have been widely adopted in the last years, and the scientific workflow management systems have become strong candidates for being migrated to the Cloud. In this research work we aim at empirically evaluating multiple Cloud providers and their corresponding optimized and non-optimized IaaS offerings with respect to their offered performance, and its impact on the incurred monetary costs when migrating and executing a workflow-based simulation environment. The experiments show significant performance improvements and reduced monetary costs when executing the simulation environment in off-premise Clouds.

Short Papers

Paper Nr: 14

Title: Executing Bag of Distributed Tasks on Virtually Unlimited Cloud Resources

Authors: Long Thai, Blesson Varghese and Adam Barker

Bag-of-Distributed-Tasks (BoDT) application is the collection of identical and independent tasks each of which requires a piece of input data located around the world. As a result, Cloud computing offers an effective way to execute BoT application as it not only consists of multiple geographically distributed data centres but also allows a user to pay for what is actually used. In this paper, BoDT on the Cloud using virtually unlimited cloud resources is investigated. To this end, a heuristic algorithm is proposed to find an execution plan that takes budget constraints into account. Compared with other approaches, for the same given budget, the proposed algorithm is able to reduce the overall execution time up to 50%.

Paper Nr: 17

Title: Automatic Abstraction of Flow of Control in a System of Distributed Software Components

Authors: Nima Kaviani, Michael Maximilien, Ignacio Silva-Lepe and Isabelle Rouvellou

CloudFoundry (CF) provides an open source platform-as-a-service software for deploying scalable software systems to the cloud. The architecture for CF is distributed by design and consists of several components which interact with one another through a message-oriented middleware. This message-oriented distributed design delivers on the scalability and resiliency requirements of the platform. In such a complex distributed multicomponent system, there is a steep learning curve for software developers to understand how components interact, what messages are exchanged between them, and how the message exchanges affect the behaviour of the system. In particular developers find it difficult to identify the execution flows, the authentication flows, interactions with the persistence layer, etc. We have developed a framework that allows interpreting the behaviour of the system by analyzing the exchanged messages between components, inspecting message contents, and extracting data and control flow across components. The paramount aim is to improve developers' understandability of the system and to examine software resiliency through approaches like bug injection and message alterations. An initial version of our framework was released to the CF community and we have collected feedback that indeed show that we are achieving some of our goals.

Paper Nr: 65

Title: Automating Resources Discovery for Multiple Data Stores Cloud Applications

Authors: Rami Sellami, Michel Vadrine, Sami Bhiri and Bruno Defude

The production of huge amount of data and the emergence of cloud computing have introduced new requirements for data management. Many applications need to interact with several heterogeneous data stores depending on the type of data they have to manage: traditional data types, documents, graph data from social networks, simple key-value data, etc. Interacting with heterogeneous data models via different APIs, multi-data store applications imposes challenging tasks to their developers. Indeed, programmers have to be familiar with different APIs. In addition, developers need to master and deal with the complex processes of cloud discovery, and application deployment and execution. Moreover, the execution of join queries over heterogeneous data models cannot, currently, be achieved in a declarative way as it is used to be with mono-data store application, and therefore requires extra implementation effort. In this paper we propose a declarative approach enabling to lighten the burden of the tedious and non-standard tasks of discovering relevant cloud environment and deploying applications on them while letting developers to simply focus on specifying their storage and computing requirements. A prototype of the proposed solution has been developed and is currently used to implement use cases from the OpenPaaS project.

Paper Nr: 109

Title: MusicBeetle - Intelligent Music Royalties Collection and Distribution System

Authors: Carlos Serrão, Hélder Carvalho and Nelson Carvalho

Music industry has been completely disrupted by a range of new online digital services and social networking systems that has forever changed the way users and businesses experience and use music. This had a tremendous impact on the established music business models that had guided a dozen year-old industry. On what concerns business music users, i.e. businesses that make use of music as part of their own business model, and on the business relation they establish with author societies or their representatives, they are required to pay royalties for the use of music. These royalties need to be distributed and authors will have the opportunity to see their work rewarded properly. The proper distribution of royalties is a non-transparent and complex process. In this paper, the authors present a system, called MusicBeetle that enables the identification, collection and distribution of music royalties through the usage of decentralised system and low cost hardware devices.

Paper Nr: 128

Title: Context-aware MapReduce for Geo-distributed Big Data

Authors: Marco Cavallo, Giuseppe Di Modica, Carmelo Polito and Orazio Tomarchio

MapReduce is an effective distributed programming model used in cloud computing for large-scale data analysis applications. Hadoop, the most known and used open-source implementation of the MapReduce model, assumes that every node in a cluster has the same computing capacity and that data are local to tasks. However, in many real big data applications where data may be located in many datacenters distributed over the planet these assumptions do not hold any longer, thus affecting Hadoop performance. This paper addresses this point, by proposing a hierarchical MapReduce programming model where a toplevel scheduling system is aware of the underlying computing contexts heterogeneity. The main idea of the approach is to improve the job processing time by partitioning and redistributing the workload among geo-distributed workers: this is done by adequately monitoring the bottom-level computing and networking context.

Posters

Paper Nr: 10

Title: An Approach in the Design of Common Authentication Solution for a Multi-Platform Cloud Environment

Authors: Primož Cigoj, Borka Jerman Blažič and Tomaž Klobučar

The security provision within multi-platform cloud computing environment is still considered not to be properly solved due to different problems with technical and human-based origin. This paper presents an attempt to provide an authentication and authorization solution based on the single sign-on (SSO) approach for cloud service users and administrators in a multi-platform environment. The problem of authentication in cloud services is briefly introduced and the approach implemented for cloud environment with two different proprietary (VMware) and open source (OpenStack) platforms is described.

Paper Nr: 39

Title: Towards Cross-layer Monitoring of Cloud Workflows

Authors: Eric Kübler and Mirjam Minor

prospectively cloud management requires sophisticated monitoring capabilities. In this paper, we introduce a novel monitoring framework for cloud-based workflow systems called cWorkload. cWorkload integrates monitoring information from different layers of the cloud architecture. The paper puts its focus on the two-layer monitoring regarding the workflow layer and the PaaS layer. We present the layered monitoring architecture, an implementation of the two-layer cross-monitoring part, and an experimental evaluation with sample workflow data. Further, we discuss related work on cloud monitoring divided into one-layer, multi-layer, and cross-layer approaches. Our plans for future work on extending the implementation by further layers towards a cross-layer, prospective monitoring for prospective cloud management are described.

Abstract:

Area 4 - Cloud Computing Enabling Technology

Full Papers

Paper Nr: 31

Title: Secure Keyword Search over Data Archives in the Cloud - Performance and Security Aspects of Searchable Encryption

Authors: Christian Neuhaus, Frank Feinbube, Daniel Janusz and Andreas Polze

Encryption can protect the confidentiality of data stored in the cloud, but also prevents search. To solve this problem, searchable encryption schemes have been proposed that allow keyword search over document collections. To investigate the practical value of such schemes and the tradeoff between security, functionality and performance, we integrate a prototypical implementation of a searchable encryption scheme into a document-oriented database. We give an overview of the performance benchmarking results of the approach and analyze the threats to data confidentiality and corresponding countermeasures.

Abstract:

Paper Nr: 38

Title: A Many-objective Optimization Framework for Virtualized Datacenters

Authors: Fabio López Pires and Benjamín Barán

The process of selecting which virtual machines should be located (i.e. executed) at each physical machine of a datacenter is commonly known as Virtual Machine Placement (VMP). This work presents a general many-objective optimization framework that is able to consider as many objective functions as needed when solving the VMP problem in a pure multi-objective context. As an example of utilization of the proposed framework, for the first time a formulation of the many-objective VMP problem (MaVMP) is proposed, considering the simultaneous optimization of the following five objective functions: (1) power consumption, (2) network traffic, (3) economical revenue, (4) quality of service and (5) network load balancing. To solve the formulated many-objective VMP problem, an interactive memetic algorithm is proposed. Simulations prove the correctness of the proposed algorithm and its effectiveness converging to a treatable number of solutions in different experimental scenarios.

Abstract:

Paper Nr: 64

Title: CloudMPL: A Domain Specific Language for Describing Management Policies for an Autonomic Cloud Infrastructure

Authors: Marwah M. Alansari, Andre Almeida, Nelly Bencomo and Behzad Bordbar

To benefit from the advantages that Cloud Computing brings to the IT industry, management policies must be implemented as a part of the operation of the Cloud. Among others, for example, the specification of policies can be used for the management of energy to reduce the cost of running the IT system or also for security policies while handling privacy issues of users. As cloud platforms are large, manual enforcement of policies is not scalable. Hence, autonomic approaches for management policies have recently received a considerable attention. These approaches allow specification of rules that are executed via rule-engines. The process of rules creation starts by the interpretation of the policies drafted by high-rank managers. Then, technical IT staff translate such policies to operational activities to implement them. Such process can start from a textual declarative description and after numerous steps terminates in a set of rules to be executed on a rule engine. To simplify the steps and to bridge the considerable gap between the declarative policies and executable rules, we propose a domain-specific language called CloudMPL. We also design a method of automated transformation of the rules captured in CloudMPL to the popular rule-engine Drools. As the policies are changed over time, code generation will reduce the time required for the implementation of the policies. In addition, using a declarative language for writing the specifications is expected to make the authoring of rules easier. We demonstrate the use of the CloudMPL language into a running example extracted from a management energy consumption case study.

Abstract:

Paper Nr: 86

Title: Dynamic Testing and Deployment of a Contract Monitoring Service

Authors: Ellis Solaiman, Ioannis Sfyrakis and Carlos Molina-Jimenez

Internet and cloud based services involve electronic interactions that are normally regulated using service agreements (SA). Once an agreement between business partners is in place, a service can be monitored and/or enforced using an SA equivalent electronic contract. Because of the dynamic nature of such Internet and cloud based relationships, the rapidity at which electronic contracts are constructed, verified for correctness, tested, and deployed is an extremely important factor. This paper describes a model checker based framework for supporting the automated testing and deployment of electronic contracts. The central components of the framework are a contract monitoring service called the Contract Compliance Checker (CCC), the SPIN model checker, and EPROMELA, a language developed specifically for modeling electronic contracts. We describe how SPIN can be used to automatically generate execution sequences from an EPROMELA model of a contract, and how such sequences can then be used to test the correctness of the model equivalent electronic contract deployed to the CCC.

Abstract:

Paper Nr: 95

Title: ANY2API – Automated APIfication - Generating APIs for Executables to Ease their Integration and Orchestration for Cloud Application Deployment Automation

Authors: Johannes Wettinger, Uwe Breitenbücher and Frank Leymann

APIs are a popular means to expose functionality provided by Cloud-based systems, which are utilized to integrate and orchestrate application as well as management functionality in a programmatic manner. In the domain of application management, they are used to fully automate management processes, for example, to deploy Cloud-based Web applications or back-ends for mobile apps. However, as not all required functionality is exposed through an API natively, such processes additionally involve a multitude of other heterogeneous technologies such as scripting languages and deployment automation tooling. Consequently, combining different technologies in an efficient manner is a complex integration challenge. In this paper, we present a generic approach for automatically generating API implementations for arbitrary executables such as scripts and compiled programs, which are not natively exposed as APIs. This APIfication tackles the aforementioned integration challenges by unifying the invocation of heterogeneous technologies while avoiding the costly and manual wrapping of existing executables because it does not scale. We further present the modular and extensible open-source framework ANY2API that implements our APIfication approach. Furthermore, we evaluate the approach and the framework by measuring the overhead of generating and using API implementations. In addition, we conduct a detailed case study to confirm the technical feasibility of the approach.

Abstract:

Title: A Modelling Concept to Integrate Declarative and Imperative Cloud Application Provisioning Technologies**Authors:** Uwe Breitenbücher, Tobias Binz, Oliver Kopp, Frank Leymann and Johannes Wettinger

Efficient application provisioning is one of the most important issues in Cloud Computing today. For that purpose, various provisioning automation technologies have been developed that can be generally categorized into two different flavors: (i) declarative approaches are based on describing the desired goals whereas (ii) imperative approaches are used to describe explicit sequences of low-level tasks. Since modern Cloud-based business applications become more and more complex, employ a plethora of heterogeneous components and services that must be wired, and require complex configurations, the two kinds of technologies have to be integrated to model the provisioning of such applications. In this paper, we present a process modelling concept that enables the seamless integration of imperative and declarative provisioning models and their technologies while preserving the strengths of both flavors. We validate the technical feasibility of the approach by applying the concept to the workflow language BPEL and evaluate its features by several criteria.

Short Papers

Paper Nr: 22**Title:** A Hedonic Price Index for Cloud Computing Services**Authors:** Persefoni Mitropoulou, Evangelia Filiopoulou, Stavroula Tsaroucha, Christos Michalakelis and Mara Nikolaidou

Cloud computing is an innovative business model, being developed at a fast pace during the last years, offering many operational and economic benefits to both the demand and the supply side of the ICT market. Infrastructure as a Service (IaaS), which includes control of fundamental computing resources, is expected to be the fastest growing model of public cloud computing. Due to the existence of several IaaS cloud providers, there is increased competition among cloud companies, which develop different pricing models in order to meet the market demand. As a consequence, prices for cloud services are a result of a multidimensional function, shaped by the service's characteristics. The development of a suitable pricing method, based on an appropriate price index able to capture the market dynamics, is an obvious necessity. The aim of this paper is the construction of such a price index, for the IaaS model, using data from a wide range of cloud providers and a large number of price bundles. The hedonic pricing method is used to decompose cloud computing services into their constituent characteristics, obtaining estimates of the contributory value of each resource. According to the results, RAM size, CPU power and subscription turned out to be the most influential factors that shape IaaS pricing.

Paper Nr: 59**Title:** New Approach to Partitioning Confidential Resources in Hybrid Clouds**Authors:** Kaouther Samet, Samir Moalla and Mahdi Khemakhem

Today, companies use more cloud environments such as hybrid clouds. Indeed, hybrid clouds give the opportunity to better manage resources mostly when companies have no space to store more resources in their private clouds. The best solution here is to allocate the required space in public cloud at a low cost. But how can resources be partitioned in hybrid clouds while assuring confidentiality of resources moved to public cloud. Many works have been done in this context. They suppose that confidentiality is assured by using encryption methods. But with this solution the cloud provider can access the resources stored on the cloud, which weakens the confidentiality of these. This work proposes an approach to the Confidential Resources Partitioning Problem in Hybrid Clouds (CRPHC) which aims at ensuring the confidentiality of resources by grouping as much as possible the most confidential resources in private cloud and resources with low degrees of confidentiality in public cloud while minimizing the size of resources to host in public cloud and consequently reducing the storage cost. This solution allows the possibility of using non-performing encryption methods which have a reduced treatment cost compared to efficient methods. Experimentally, our solution will be evaluated and compared to optimal solution given by CPLEX.

Paper Nr: 78**Title:** Cloud Spreadsheets Supporting Data Processing in the Encrypted Domain**Authors:** D. A. Rodríguez-Silva, L. Adkinson-Orellana, B. Pedrero-López and F. J. González-Castaño

Security has become one of the main barriers for the adoption of cloud services. A range of legal initiatives that require support mechanisms such as access control and data encryption have been proposed to ensure privacy for data moved to the cloud. Although these mechanisms are currently feasible in situations in which the cloud acts as a mere data storage system, they are insufficient in more complex scenarios requiring processing in external cloud servers. Several new schemes have been proposed to overcome these shortcomings. Data Processing in the Encrypted Domain (DPED) permits arithmetic operations over ciphered data and the generation of encrypted results, without exposure of clear data. In such a set-up, the servers have no access to the information at any point of the process. In this paper we describe, as a case study of secure cloud data processing, a cloud spreadsheet that relies on DPED libraries to perform operations in the encrypted domain. Tests performed on local servers and in the Google cloud through the Google App Engine platform show that representative real applications can benefit from this technology. Because the proposed solution is PaaS-oriented, developers can apply the libraries to other applications.

Paper Nr: 91**Title:** Towards High Performance Big Data Processing by Making Use of Non-volatile Memory**Authors:** Shuichi Oikawa

Cloud computing environments for big data processing require high performance storage. There are emerging high performance memory storage technologies, such as next generation non-volatile (NV) memory and battery backed NV-DIMM. While their performance is much higher than the current block storage devices, such as SSDs and HDDs, they provide only limited capacity. Such limited capacity makes it difficult for memory storage to be adapted as mass storage, and their uses in cloud computing environments have been severely limited. This paper proposes a method that combines memory storage with block storage. It makes use of memory storage as cache of block storage in order to remove the capacity limitation of memory storage. The proposed method inherits the high performance of memory storage and also the large capacity of block storage. Therefore, memory storage can be transparently used as a part of mass storage while its low overhead access can accelerate storage performance. The proposed method was implemented as a device driver of the Linux kernel. Its performance evaluation shows that it outperforms a bare SSD drive and achieves better performance on the Hadoop and database environments.

Paper Nr: 99**Title:** The Docker Ecosystem Needs Consolidation**Authors:** René Peinl and Florian Holzschuher

Docker provides a good basis to run composite applications in the cloud, especially if those are not cloud-aware, or cloud-native. However, Docker concentrates on managing containers on one host, but SaaS providers need a container management solution for multiple hosts. Therefore, a number of tools emerged that claim to solve the problem. This paper classifies the solutions, maps them to requirements from a case study and identifies gaps and integration requirements. We conclude that the Docker ecosystem could help moving from IaaS and PaaS solutions towards a runtime environment for SaaS applications, but needs consolidation.

Paper Nr: 107

Title: Container-based Virtualization for HPC

Authors: Holger Gantikow, Sebastian Klingberg and Christoph Reich

Abstract: Experts argue that the resource demands of High Performance Computing (HPC) clusters request bare-metal installations. The performance loss of container virtualization is minimal and close to bare-metal, but in comparison has many advantages, like ease of provisioning. This paper presents the use of the newly adopted container technology and its multiple conceptual advantages for HPC, compared to traditional bare-metal installations or the use of VMs. The setup based on Docker (Docker, 2015) shows a possible use in private HPC sites or public clouds as well. The paper ends with a performance comparison of a FEA job run both bare-metal and using Docker and a detailed risk analysis of Docker installations in a multi-tenant environment, as HPC sites usually are.

Paper Nr: 119

Title: Towards Self-Protective Multi-Cloud Applications - MUSA – a Holistic Framework to Support the Security-Intelligent Lifecycle Management of Multi-Cloud Applications

Authors: Erkuden Rios, Eider Iturbe, Leire Orue-Echevarria, Massimiliano Rak and Valentina Casola

Abstract: The most challenging applications in heterogeneous cloud ecosystems are those that are able to maximise the benefits of the combination of the cloud resources in use: multi-cloud applications. They have to deal with the security of the individual components as well as with the overall application security including the communications and the data flow between the components. In this paper we present a novel approach currently in progress, the MUSA framework. The MUSA framework aims to support the security-intelligent lifecycle management of distributed applications over heterogeneous cloud resources. The framework includes security-by-design mechanisms to allow application self-protection at runtime, as well as methods and tools for the integrated security assurance in both the engineering and operation of multi-cloud applications. The MUSA framework leverages security-by-design, agile and DevOps approaches to enable the security-aware development and operation of multi-cloud applications.

Paper Nr: 120

Title: High Performance Virtual Machine Recovery in the Cloud

Authors: Valentina Salapura and Richard Harper

Abstract: In this paper, we outline and illustrate concepts that are essential to achieve fast, highly scalable virtual machine planning and failover at the Virtual Machine (VM) level in a data center containing a large number of servers, VMs, and disks. To illustrate the concepts a solution is implemented and analyzed for IBM's Cloud Managed Services enterprise cloud. The solution enables at-failover-time planning, and keeps the recovery time within tight service level agreement (SLA) allowed time budgets via parallelization of recovery activities. The initial serial failover time was reduced for an order of magnitude due to parallel VM restart, and to parallel VM restart combined with parallel storage device remapping.

Paper Nr: 127

Title: Adopting an Agent and Event Driven Approach for Enabling Mutual Auditability and Security Transparency in Cloud based Services

Authors: Moussa Ouedraogo, Eric Dubois, Djamel Khadraoui, Sebastien Poggi and Benoit Chenal

Abstract: We propose an event-driven approach for the automated audit of cloud based services security. The proposed approach is a solution to two of the intrinsic security issues of cloud based services, notably the need of security transparency and mutual auditability amongst the stakeholders. We leverage a logic based event specification language to represent patterns of events which occurrence can be evidence of security anomaly or breach or simply a sign of a nefarious use of the cloud infrastructure by some of its users. The use of dedicated algorithms for the detection of composite events coalesced with the definition of primitive events structure based on XCCDF format ensures the reuse and interoperability with security audit tools based on the Security Content and Automation Protocol-SCAP. The implementation and application of the approach on a cloud service dealing with electronic archiving have demonstrated its feasibility and viability.

Posters

Paper Nr: 80

Title: SLAFM - A Service Level Agreement Formal Model for Cloud Computing

Authors: Lucia De Marco, Filomena Ferrucci and M-Tahar Kechadi

Abstract: Cloud Computing services are regulated by a contract called Service Level Agreement (SLA). It is co-signed between the customers and the providers after a negotiation phase, and during their validity time several constraints have to be respected by the involved parties. Due to their popularity, cloud services are enormously used and unfortunately also abused, specially by cyber-criminals. Sometimes the crimes have the consequence of violating some contractual constraints without the parties are aware of. A manner for guaranteeing more control of the SLA respect is to consider a dedicated system interacting with the cloud services and detecting the SLA violations by analysing the log files. We introduce a formal model aimed to represent the contents of such SLAs necessary in the context of an automatic mechanism for detecting SLA violations.

Area 5 - Mobile Cloud Computing and Services

Full Papers

Paper Nr: 43

Title: The Case for Visualization as a Service - Mobile Cloud Gaming as an Example

Authors: Abdelmounaam Rezgui and Zaki Malik

Abstract: In recent years, significant progress has been made to improve the power efficiency of mobile devices. In particular, new GPU architectures have made it possible to run compute-intensive applications directly on battery-powered mobile devices. In parallel, research is also being conducted in the area of application offloading, the process of running compute-intensive tasks on cloud servers and delivering the results of these computations to mobile devices through their wireless interfaces. It is important to understand the power consumption implications of each of these two options. In this paper, we use mobile cloud gaming as an

example to evaluate and compare these two alternatives (running games on the cloud or on mobile devices.) Based on this comparison, we introduce the concept of Visualization as a Service (VaaS) as a new model to design and implement graphics-intensive applications for mobile devices. In this model, advanced visualization capabilities (e. g., interactive visualization of high resolution videos/images) would be provided to mobile users as a service via the Internet. We show through actual hardware specifications that, despite the recent introduction of ultra low power GPUs for mobile devices, it remains far more power efficient to offload graphics-intensive tasks to the cloud. The associated latency can still be tolerated in most applications.

Paper Nr: 79

Title: Cloud-side Execution of Database Queries for Mobile Applications

Authors: Robert Pettersen, Steffen Viken Valvåg, Åge Kvalnes and Dag Johansen

Abstract: We demonstrate a practical way to reduce latency for mobile .NET applications that interact with cloud database services. We provide a programming abstraction for location-independent code, which has the potential to execute either locally or at a satellite execution environment in the cloud, in close proximity to the database service. This preserves a programmatic style of database access, and maintains a simple deployment model, but allows applications to offload latency-sensitive code to the cloud. Our evaluation shows that this approach can significantly improve the response time for applications that execute dependent queries, and that the required cloud-side resources are modest.

Short Papers

Paper Nr: 124

Title: Telco Clouds - Modelling and Simulation

Authors: Jakub Krzywda, William Tärneberg, Per-Olov Östberg, Maria Kihl and Erik Elmroth

Abstract: In this paper, we propose a telco cloud meta-model that can be used to simulate different infrastructure configurations and explore their consequences for system performance and costs. To achieve this, we analyse current telecommunication and data centre infrastructure paradigms, describe the architecture of the telco cloud, and detail the benefits of merging both infrastructures in a unified system. Next, we detail the dynamics of the telco cloud and identify the components that are the most relevant from the perspective of modelling performance and cost. As a number of well established simulation technologies exist for most of the telco cloud components, we survey existing models in an attempt to construct a suitable composite meta-model. Finally, we present a showcase scenario to demonstrate the scope of our telco cloud simulator.
