

(<http://www.eubra-bigsea.eu>)

International Workshop On Assured Cloud Computing And QoS Aware Big Data

[HOME \(/\)](#) / [NEWS & EVENTS \(/LIST-NEWS-EVENTS\)](#) / [EVENTS \(/EVENTS\)](#)



Agenda

LOCATION: room Serrano

14:00-16:00 Session 3D: WACC 2017 (I)

CHAIRS:

Ignacio Blanquer (Universitat Politècnica de València, Spain)


Roy Campbell (University of Illinois at Urbana-Champaign, USA)

Wagner Meira Jr. (Universidade Federal do Minas Gerais, Brazil)

14:00 **Keynote: US Air Force Interests and Directions in Cyber Security**


Colonel Ryan Thomas, Deputy Commander European Office of Aerospace Research and Development (AFOSR/EOARD)

Abstract: The basic science arm of the Air Force Research Laboratory, the Air Force Office of Scientific Research (AFOSR), funds basic research into the science of security. The goal of this research portfolio is to enable the development safe, secure and dependable information systems. With the widespread application of cloud services in the commercial and military domains, many of these results have applications to cloud, big data, and distributed computing problems. This talk will discuss the cyber security and international research interests of AFOSR. Highlights of AFOSR sponsored research will be presented, including new work in areas such as interactive and automated theorem proving, behavior based access control, big data policy monitoring, and safe machine learning techniques. Additionally, opportunities will be shared for working on research with the US Air Force through collaboration, travel, and funded research.

- 14:30 **IT Security and Privacy Standards in Comparison Improving FedRAMP Authorization for Cloud Service Providers**  (http://www.eubra-bigsea.eu/sites/default/files/WACC2017_IT%20Security%20and%20Privacy%20Standards%20in%20Comparison.pdf)


Carlo Di Giulio (University of Illinois at Urbana-Champaign, USA)
 Read Sprabery (University of Illinois at Urbana-Champaign, USA)
 Charles Kamhoua (Air Force Research Laboratory, USA)
 Kevin Kwiat (Air Force Research Laboratory, USA)
 Roy Campbell (University of Illinois at Urbana-Champaign, USA)
 Masooda Bashir (University of Illinois at Urbana-Champaign, USA)

Abstract: To demonstrate compliance with privacy and security principles, information technology (IT) service providers often rely on security standards and certifications. However, the appearance of new service models such as cloud computing has brought new threats to information assurance, weakening the protection that existing standards can provide. In this study, we analyze four highly regarded IT security standards used to assess, improve, and demonstrate information systems assurance and cloud security. ISO/IEC 27001, SOC 2, C5, and FedRAMP are standards adopted worldwide and constantly updated and improved since the first release of ISO in 2005. We examine their adequacy in addressing current threats to cloud security, and provide an overview of the evolution over the years of their ability to cope with threats and vulnerabilities. By comparing the standards alongside each other, we investigate their complementarity, their redundancies, and the level of protection they offer to information stored in cloud systems. We unveil vulnerabilities left unaddressed in the four frameworks, thus questioning the necessity of multiple standards to assess cloud assurance. We suggest necessary improvements to meet the security requirements made indispensable by the current threat landscape.

- 14:50 **A Game-Theoretic Approach for Runtime Capacity Allocation in MapReduce**  (http://www.eubra-bigsea.eu/sites/default/files/WACC2017_A%20Game-Theoretic%20Approach%20for%20Runtime%20Capacity%20Allocation%20in%20MapReduce.pdf)


Eugenio Gianniti (Politecnico di Milano, Italy)
 Danilo Ardagna (Politecnico di Milano, Italy)
 Michele Ciavotta (Politecnico di Milano, Italy)
 Mauro Passacantando (Università di Pisa, Italy)

Abstract: Nowadays many companies have available large amounts of raw, unstructured data. Among Big Data enabling technologies, a central place is held by the MapReduce framework and, in particular, by its open source implementation, Apache Hadoop. For cost effectiveness considerations, a common approach entails sharing server clusters among multiple users. The underlying infrastructure should provide every user with a fair share of computational resources, ensuring that service level agreements (SLAs) are met and avoiding wastes. In this paper we consider mathematical models for the optimal allocation of computational resources in a Hadoop 2.x cluster with the aim to develop new capacity allocation techniques that guarantee better performance in shared data centers. Our goal is to get a substantial reduction of power consumption while respecting the deadlines stated in the SLAs and avoiding penalties associated with job rejections. The core of this approach is a distributed algorithm for runtime capacity allocation, based on Game Theory models and techniques, that mimics the Map- Reduce dynamics by means of interacting players, namely the central Resource Manager and Class Managers.

- 15:10 **Automatic Consolidation of Virtual Machines in On-Premises Cloud Platforms**  (http://www.eubra-bigsea.eu/sites/default/files/WACC2017_%20Virtual%20Machine%20Consolidation%20Agent%20VMCA.pdf)

Carlos De Alfonso (Universidad Politécnica de Valencia, Spain)
 Ignacio Blanquer (UPV, Spain)
 Germán Moltó (Universitat Politècnica de València, Spain)
 Miguel Caballer (Universitat Politècnica de València, Spain)

Abstract: After a sequence of creation and destruction of virtual machines (VMs) in an on-premises Cloud computing platform, the scheduling decisions to host the VMs are far from being optimal and the fragmentation of the physical resources may impede the platform to host some VMs despite the free available virtualization resources. This paper describes a Virtual Machine Consolidation Agent that addresses this problem by analyzing the distribution of the VMs in the virtualization platform to migrate some of them among hosts, in order to defragment the physical resources and to enhance the efficiency on their usage. The agent has been validated in a production platform, where it is capable of minimizing the number of servers needed to host the VMs. The algorithms achieve near-optimal values at a very reduced computational cost, thus making it suitable for production platforms.

- 15:30 **PRIVAaaS: privacy approach for a distributed cloud-based data analytics platforms**  (http://www.eubra-bigsea.eu/sites/default/files/WACC2017_PRIVAAaaS%20Privacy%20Approach%20for%20a%20Distributed%20cloud-based%20data%20analytics%20platforms.pdf)

Tania Basso (UNICAMP, Brazil)
 Regina Moraes (UNICAMP, Brazil)
 Nuno Antunes (University of Coimbra, Portugal)
 Marco Vieira (University of Coimbra, Portugal)
 Walter Santos (Federal University of Minas Gerais, Brazil)
 Wagner Meira (Federal University of Minas Gerais, Brazil)

Abstract: Assuring data privacy is a key challenge that is exacerbated by Big Data storage and analytics processing requirements. Big Data and Cloud Computing are inseparable allowing the users to access data from any device, making data privacy essential as the data sets are exposed through the web. Organizations care about data privacy as it directly affects the confidence that clients have that their personal data are safe. This paper presents a data privacy approach - PRIVAaaS, which was integrated to the LEMONADE Web-based platform, developed to compose ETL and Machine Learning workflows. The 3-level approach of PRIVAaaS, based on data anonymization policies, is implemented in a software toolkit that provides a set of libraries and tools which allows controlling and reducing data leakage in the context of big data processing.

16:30-18:00 Session 4D: WACC 2017 (II)

CHAIRS:



Ignacio Blanquer (Universitat Politècnica de València, Spain)
 Roy Campbell (University of Illinois at Urbana-Champaign, USA)
 Wagner Meira Jr. (Universidade Federal do Minas Gerais, Brazil)

16:30 **Keynote: Assured Cloud Computing: are we there yet?**

Paulo Esteves-Veríssimo, University of Luxembourg FSTC - Faculty of Science, Technology, and Communication - CSC SnT - Interdisciplinary Centre for Security, Reliability and Trust

17:00 **A lightweight MapReduce framework for secure processing with SGX** (http://www.eubra-bigsea.eu/sites/default/files/WACC2017_%20A%20lightweight%20MapReduce%20framework%20for%20secure%20processing%20with%20SGX.pdf)

Rafael Pires (University of Neuchatel, Switzerland)
 Daniel Gavril (Alexandru Ioan Cuza University of Iasi, Romania)
 Pascal Felber (University of Neuchatel, Switzerland)
 Emanuel Onica (Alexandru Ioan Cuza University of Iasi, Romania)
 Marcelo Pasin (Université de Neuchâtel, Switzerland)

Abstract: MapReduce is a programming model used extensively for parallel data processing in distributed environments. A wide range of algorithms were implemented using MapReduce, from simple tasks like sorting and searching up to complex clustering and machine learning operations. Many of these implementations are part of services externalized to cloud infrastructures. Over the past years, however, many concerns have been raised regarding the security guarantees offered in such environments. Some solutions relying on cryptography were proposed for countering threats but these typically imply a high computational overhead. Intel, the largest manufacturer of commodity CPUs, recently introduced SGX (software guard extensions), a set of hardware instructions that support execution of code in an isolated secure environment. In this paper, we explore the use of Intel SGX for providing privacy guarantees for MapReduce operations, and based on our evaluation we conclude that it represents a viable alternative to a cryptographic mechanism. We present results based on the widely used k-means clustering algorithm, but our implementation can be generalized to other applications that can be expressed using MapReduce model.

17:20 **Evaluating the performance of continuous test-based cloud service certification** (http://www.eubra-bigsea.eu/sites/default/files/WACC2017_Evaluating%20the%20performance%20of%20continuous%20test-based%20cloud%20service%20certification.pdf)

Philipp Stephanow (Fraunhofer AISEC, Germany)
 Christian Banse (Fraunhofer AISEC, Germany)

Abstract: Continuous test-based cloud certification uses tests to automatically and repeatedly evaluate whether a cloud service satisfies customer requirements over time. However, inaccurate tests can decrease customers' trust in test results and can lead to providers disputing results of test-based certification techniques. In this paper, we propose an approach how to evaluate the performance of test-based cloud certification techniques. Our method allows to infer conclusions about the general performance of test-based techniques, compare alternative techniques, and compare alternative configurations of test-based techniques. We present experimental results on how we used our approach to evaluate and compare exemplary test-based techniques which support the certification of requirements related to security, reliability and availability.

17:40 **A game theoretic method for VM-to-hypervisor attacks detection in cloud environment**

Amin Nezarat (PNU University, Iran)

Abstract: Cloud computing is a pool of scalable virtual resources serving a large number of users who pay fees depending on the extent of utilized service. From payment perspective, cloud is like electricity and water as people who use more of this shared pool should pay larger fees. Cloud computing involves a diverse set of technologies including networking, virtualization, transaction scheduling, etc. so it is vulnerable to a wide range of security threats. Some of the most important security issues threatening the cloud computing systems originate from virtualization technology, as it constitutes the main body and basis of these systems. The most important virtualization-based security threats include VM side-channel, VM Escape and Rootkit attacks. The previous works on the subject of virtualization security rely on hardware approaches such as the use of firewalls, which are expensive, the use of schedulers to control the side channels along with noise injection, which impose high overhead, or the use of agents to collect information and send them back to a central intrusion detection system, which itself can become the target of attacker. In the method presented in this paper, a group of mobile agents act as the sensors of invalid actions in the cloud environment. They start a non-cooperative game with the suspected attacker, and then calculate the Nash equilibrium value and utility so as to differentiate an attack from legitimate requests and determine the severity of attack and its point of origin. The simulation results show that this method can detect the attacks with 86% accuracy. The use of mobile agents and their trainability feature has led to reduced system overhead and accelerated detection process.

Scope

WACC draws together researchers, practitioners, and thought leaders from government, industry, and academia. The workshop provides a forum of dialogue centered upon the development and advancement of an effort to design, implement, and evaluate dependable cloud architectures that can provide assurances with respect to security, reliability, and timeliness of computations (or services). Some new "assured" target applications include, but are not limited to, dependable Big Data applications and streaming, data analytics and its tools, real-time computations for monitoring, control of cyber-physical systems such as power systems, and mission critical computations for rescue and recovery.

The technical emphasis of WACC is design, implementation, and evaluation of cloud services, data analytics tools, and security solutions to enable dependable Big Data applications. Research on cloud services, ICT-skilled data scientists and application developers can find complementary solutions and partnerships to evaluate and integrate additional solutions. Data scientists can find new tools that could address existing needs.

Topics

The *International Workshop on Assured Cloud Computing and QoS aware Big Data* (WACC 2017) is held in conjunction with the *17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing* in Madrid, Spain. The workshop will include a mix invited talks, refereed papers, and panels. Papers will appear in the workshop proceedings to be published by the IEEE and CCGrid.

Topics of Interest

We invite submissions on any topic related to cloud computing that aligns with the workshop goals listed above.

- Assurance in clouds and their applications including time-critical cloud computations, streaming and the QoS of cloud services.
- Assurances and early validation for Big Data, its algorithms, innovative data models, analytic solutions, and applications.

- Experiences on the proactive and reactive allocation of cloud resources for Big Data applications.
- Dependability challenges in cloud environments and novel techniques including dependable performance for Big Data applications.
- Security and availability of the networks for assured cloud computing.
- State of the art techniques to provide cloud resiliency, cloud resource management for Big Data, on-demand elasticity for Big Data, reliable and scalable cloud services.
- Role of virtualization and multicore hardware in assuring security and reliability.
- Validation using experimental and/or formal methods.
- Integrity verification and attestation.
- Metrics and risk models for assessing assurance of cloud environments.
- Game theory and mathematical modelling of cloud computing security.
- Assured services, protocols and standards for clouds.
- Secure data management and computation outsourcing.
- Security and privacy aspects in cloud services, efficient data protection mechanisms for Big Data and data privacy at storage layer.
- Usability of assurance provisions, interoperability, use cases.
- Cost of assured clouds, energy, trade-offs, budget.
- Identity management and authentication.
- Policy management access controls, authorization for clouds and Big Data.
- Trust and Regulatory Compliance.

Organisers

Workshop Organisers

Ignacio Blanquer, Universitat Politècnica de València, ES
 Rakesh Bobba, Oregon State University, US
 Andrey Brito, Federal University of Campina Grande, BR
 Roy Campbell, University of Illinois at Urbana-Champaign, US
 Christof Fetzer, Dresden University of Technology, DE
 Zbigniew Kalbarczyk, University of Illinois at Urbana-Champaign, US
 Charles Kamhoua, Air Force Research Laboratory, US
 Kevin Kwiat, Air Force Research Laboratory, US
 Wagner Meira, Jr., Federal University of Minas Gerais, BR
 Luigi Romano, Synclab S.r.l., IT

Technical Program Committee

Jussara Almeida, Federal University of Minas Gerais, BR
 Nazareno Andrade, Federal University of Minas Gerais, BR
 Nuno Antunes, University of Coimbra, PT
 Rosa Badia, Barcelona Supercomputing Center, ES
 Enrico Barbierato, Politecnico di Milano, IT
 Tania Basso, University of Campinas, BR
 Walter dos Santos, Federal University of Minas Gerais, BR
 Pascal Felber, University of Neuchatel, CH
 Walter Filho, Federal University of Minas Gerais, BR
 Sandro Flore, CMCC, IT
 Karthik Gopalan, Binghamton University, US
 Dorgival Guedes, Federal University of Minas Gerais, BR
 Yier Jin, University of Central Florida, US
 Stefan Kopsell, Dresden University of Technology, DE
 Daniele Lezzi, Barcelona Supercomputing Center, ES
 Chen Liu, Clarkson University, US
 Andrew Martin, Oxford University, UK
 Andre Martin, Dresden University of Technology, DE
 German Moltó, Universitat Politècnica de València, ES
 Laurent Njilla, Air Force Research Laboratory, US
 Karthik Pattabiraman, University of British Columbia, CA
 Peter Pietzuch, Imperial College, UK
 Nádia Puchalski Kozievitch, Federal Technical University of Parana, BR
 Harigovind V. Ramasamy, IBM, US
 Sachin Shetty, Old Dominion University, US
 Jian Tang, Syracuse University, US
 Aad van Moorsel, University of Newcastle, UK
 Marco Viera, University of Coimbra, PT
 Monica Vitali, Politecnico di Milano, IT
 Long Wang, IBM, US
 Lok Yan, Air Force Research Laboratory, US
 Ming Zhao, Arizona State University, US

Registration & Logistics

Registration

Please view the main conference registration page for details <https://www.arcos.inf.uc3m.es/wp/ccgrid2017/for-attendees/>
 (<https://www.arcos.inf.uc3m.es/wp/ccgrid2017/for-attendees/>)

Logistics

The conference will be held in Madrid, at Melia Los Galgos Hotel (<http://www.melia.com/en/hotels/spain/madrid/melia-galgos/index.html>):

Calle de Claudio Coello, 139
Madrid 28006
Spain

For suggestions regarding accommodation visit <https://www.arcos.inf.uc3m.es/wp/ccgrid2017/for-attendees/hotel-reservat...>
(<https://www.arcos.inf.uc3m.es/wp/ccgrid2017/for-attendees/hotel-reservation/>)



EUBra-BIGSEA is funded by the European Commission under the Cooperation Programme, Horizon 2020 grant agreement No 690116. Este projeto é resultante da 3a Chamada Coordenada BR-UE em Tecnologias da Informação e Comunicação (TIC), anunciada pelo Ministério de Ciência, Tecnologia e Inovação (MCTI). | **Disclaimer** (<http://www.eubra-bigsea.eu/disclaimer-privacy-policy>) | **Privacy Policy** | (<http://www.eubra-bigsea.eu/content/privacy-policy-summary>)