Guest Editors' Introduction: Securing Big Data Applications in the Cloud

Bharat Bhargava, Purdue University
Ibrahim Khalil, RMIT University, Australia
Ravi Sandhu, University of Texas, San Antonio

Traditional security mechanisms are tailored for small-scale data, so they don't meet the needs of big data analytics and storage applications. This special issue aims to stimulate discussion and research toward the innovation of security and privacy mechanisms for big data applications in a cloud environment.

loud-based platforms are playing an increasingly important role in the context of big data analytics and storage applications. The velocity, volume, and variety of big data for large-scale cloud infrastructures can't be enhanced without security and privacy. Because traditional security mechanisms are tailored to securing small-scale data, they can't meet the needs of big data. Moreover, the inherent vulnerabilities of a cloud-based environment require significant focus on both privacy and security together with risk management procedures. To stimulate discussion and invigorate research interest toward the innovation of security and privacy mechanisms for big data applications in a cloud environment, this special issue discusses topics such as intrusion detection and attack prevention, risk awareness, secure and efficient data sharing, and access control.

The call for papers was well timed given the dynamic ongoing research on security and privacy for cloud-based big data applications. We received numerous submissions, and, after a rigorous peer review process, we selected five articles for this special issue.

The Articles

In "Enhancing Big Data Security with Collaborative Intrusion Detection," Zhiyuan Tan and his colleagues introduce a collaborative intrusion detection framework that focuses on efficiency, scalability, and self-adaption for big data applications in cloud computing. The system performs intrusion detection at both the host and network levels in a collaborative manner, using a model for parallel network summarization that utilizes cloud computing features.

The article, "Risk-Aware Virtual Resource Management for Multitenant Cloud Datacenters," by Abdulrahman A. Almutairi and Arif Ghafoor, presents efficient risk-aware virtual resource management procedures that avoid information leakage in cloud-based multitenant sharing environments. The authors propose a sharing-based heuristic that reduces overall risk, and a partition-based heuristic that is scalable for large datacenters. They use sensitivity characterization to address the virtual resource as-

signment problem in environments with role-based access control (RBAC).

In "Efficient and Secure Transfer, Synchronization, and Sharing of Big Data," Kyle Chard, Steven Tuecke, and Ian Foster propose secure and efficient data access, transfer, and sharing functions for large datasets across multiple types of local and cloud storage, which they achieve through the Globus software-as-a-service (SaaS) platform for data transfer and synchronization. Their secure framework supports resiliency and integrity while spanning a variety of heterogeneous data storage systems.

A fourth article, "Location-Based Security Framework for Cloud Perimeters," by Chetan Jaiswal, Mahesh Nath, and Vijay Kumar, proposes a cost-effective model for location-based firewall filtering of attacks for mobile and static cloud environments. The authors introduce two schemes for identifying and filtering out static and mobile security attackers using a logic-based framework that's coupled with the dynamic revision of firewall policies. These functions are performed in a distributed manner, keeping the local and global policies in sync.

Finally, in "Multilabels-Based Scalable Access Control for Big Data Applications," Chen Hongsong, Bharat Bhargava, and Fu Zhongchuan propose a multilabel-based access control approach for Hadoop-based big data applications in clouds that is both efficient and scalable. The work combines active bundle, RBAC, discretionary access control (DAC), and mandatory access control (MAC), and includes a security degree, lifetime, and access policy among the multilabels. The authors evaluate the approach using a rigorous case study of a personal health record (PHR) data storage application. As both coauthor and guest editor, Bharat Bhargava did not take part in the peer review of this article.

We thank all of the authors who submitted manuscripts to this special issue. We also wish to thank the reviewers who helped to review the papers in a very short time period, as well as Editor in Chief Mazin Yousif for his encouragement and support in organizing this special issue. Finally, we thank the publication staff for their continuous

support. We close this editorial by noting that several more feature topics on scalable and secure big data analytics are due to appear in the magazine in the near future. ●●●

BHARAT BHARGAVA *is a professor of computer science at Purdue University. His research interests include security and privacy issues in distributed systems and sensor networks. This involves identity management, secure routing and dealing with malicious hosts, adaptability to attacks, and experimental studies. His recent work involves attack graphs for collaborative attacks. Bhargava has a PhD in computer science from Rutgers University. Contact him at bbshail@purdue.edu.*

IBRAHIM KHALIL *is a senior lecturer in the School of Computer Science and IT, RMIT University, Melbourne, Australia. His research interests include data clustering, network security, scalable computing in distributed systems, m-health, e-health, wireless and body sensor networks, biomedical signal processing, and remote healthcare. Khalil has a PhD in computer science from the University of Berne, Switzerland. Contact him at Ibrahim.khalil@rmit.edu.au.*

RAVI SANDHU *is the executive director of the Institute for Cyber Security at the University of Texas, San Antonio, where he holds the Lutcher Brown Endowed Chair in Cyber Security. His research interests include cybersecurity practice and education. Sandhu has a PhD in computer science from Rutgers University. He is an IEEE, ACM, and AAAS Fellow. Contact him at ravi.sandhu@utsa.edu.*