Searching over Encrypted Data in Cloud Systems

Florian Kerschbaum
SAP Research
Karlsruhe, Germany
florian.kerschbaum@sap.com

ABSTRACT

Security is still a major inhibitor of cloud computing. When companies are testing cloud applications, e.g. for storage or databases, they use generated data for fear of data loss. Modern encrypted databases where the cryptographic key remains at the client provide a solution to this problem. Recent results in cryptography, such order-preserving encryption, and database systems [3] enable the practical use of these systems. We report on our pre-development efforts of implementing such an encrypted database in an in-memory, column store database [1]. We highlight some unsolved research challenges: such as access control, infrequent queries and security vs. performance query optimization. Challenges to key management in multi-user environments remain largely unsolved [2]. We give an overview of the architecture and performance benchmarks on our prototype which are very encouraging for practical adoption.

The talk is structured in three parts:

1. We will give *background* on the architecture of the cloud database. First, we present overviews of recent developments in cryptography, e.g., order-preserving encryption, searchable encryption, proxy re-encryption and somewhat homomorphic encryption. Second, we give an introduction to new in-memory, column-store database architectures including compression techniques such as order-preserving dictionaries and multi-core database operators.

2. We highlight some *research challenges*, such as access control, infrequent queries, security vs. performance trade-offs and key management. Particularly, in multi-user environments databases need to handle access control, as well as sophisticated key management. Furthermore, when using adjustable (onion) encryption selection queries can modify the (security) state of the database.

3. We report on some *initial pre-development results* in our research group. We implemented a prototype and can give some performance figures. Furthermore, we show our progress on some of the outlined challenges.

SACMAT'13, June 12–14, 2013, Amsterdam, The Netherlands.
ACM 978-1-4503-1950-8/13/06.

Categories and Subject Descriptors

C.2.4 [**Computer-Communication Networks**]: Distributed Systems—*Distributed databases*; D.4.6 [**Operating Systems**]: Security and Protection—*Cryptographic controls*

Keywords

Cloud, Database, Encryption

Short Biography

Florian Kerschbaum is a chief expert in the security research department at SAP in Karlsruhe, Germany. In the academic year 2011/12 he was on leave as the deputy professor for the chair of privacy and data security at Dresden University of Technology. His research interests center around security and privacy algorithms and protocols for the next-generation, cross-organizational business applications. He holds a Ph.D. in computer science from the Karlsruhe Institute of Technology, a master's degree from Purdue University, and a bachelor's degree from Berufsakademie Mannheim.

Acknowledgements

I would like to thank my colleagues Andreas Schaad, Axel Schröpfer and Mathias Kohler, Martin Härterich, Walter Tighzert, Patrick Grofig working with me on this exciting project. Also, I would like to give a special thanks to our former manager Stephan Fischer for initiating the project.

REFERENCES

[1] S. Hildenbrand, D. Kossmann, T. Sanamrad, C. Binnig, F. Färber, and J. Wöhler. Query processing on encrypted data in the cloud. *Department of Computer Science, Technical Report 735, ETH Zürich*, 2011.

[2] F. Kerschbaum. Collusion-resistant outsourcing of private set intersection. In *Proceedings of the 27th ACM Symposium On Applied Computing (SAC)*, 2012.

[3] R. Popa, C. Redfield, N. Zeldovich, and H. Balakrishnan. CryptDB: Protecting confidentiality with encrypted query processing. In *Proceedings of the 23rd ACM Symposium on Operating Systems Principles (SOSP)*, 2011.