Cryptology ePrint Archive: Report 2011/244

PRISM -- Privacy-Preserving Search in MapReduce

Erik-Oliver Blass and Roberto Di Pietro and Refik Molva and Melek Onen

Abstract: We present PRISM, a privacy-preserving scheme for word search in cloud computing. Assuming a curious cloud provider, privacy of data stored in the cloud becomes an issue. The main challenge in the context of cloud computing is to design a scheme that achieves privacy while preserving the efficiency of cloud computing. Solutions from related research, like encrypted keyword search or Private Information Retrieval (PIR), fall short of meeting real-world cloud requirements and are impractical. PRISM's idea is to transform the problem of word search into a set of parallel instances of PIR on small datasets. Each PIR instance on a small dataset is efficiently solved by a node in the cloud during the ``Map'' phase of MapReduce. Outcomes of map computations are then aggregated during the ``Reduce'' phase. Due to the linearity of PRISM, the simple aggregation of map results yields the final output of the word search operation. We have implemented PRISM on Hadoop MapReduce and evaluated its efficiency using real-world DNS logs. The overhead of PRISM over non-private search is only 11%. Thus, PRISM offers privacy-preserving search that meets cloud computing efficiency requirements. Moreover, PRISM is compatible with standard MapReduce, not requiring any change to the interface or infrastructure.

Category / Keywords: Cloud computing, security, privacy, word search, MapReduce

Date: received 16 May 2011, last revised 21 Feb 2012

Contact author: blass at ccs neu edu

Available formats: PDF | BibTeX Citation

Version: 20120221:185418 (All versions of this report)

Discussion forum: Show discussion | Start new discussion

[Cryptology ePrint archive]