'Melbourne Shuffle' secures data in the cloud

10 July 2014

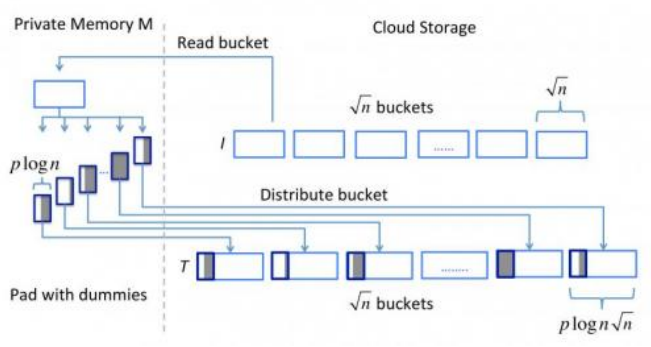

Encryption might not be enough for all that data stored in the cloud. An analysis of usage patterns -- which files are accessed and when -- can give away secrets as well. Computer scientists at Brown have developed an algorithm to sweep away those digital footprints. It's a complicated series of dance-like moves they call the Melbourne Shuffle. Credit: Tamassia Lab / Brown University

To keep data safe in the cloud, a group of computer scientists suggests doing the Melbourne Shuffle.

That may sound like a dance move (and it is), but it's also a computer algorithm developed by researchers at Brown University.

The computing version of the Melbourne Shuffle aims to hide patterns that may emerge as users access data on cloud servers. Patterns of access could provide important information about a dataset—information that users don't necessarily want others to know—even if the data files themselves are encrypted.

"Encrypting data is an important security measure. However, privacy leaks can occur even when accessing encrypted data," said Olga Ohrimenko, lead author of a paper describing the algorithm. "The objective of our work is to provide a higher level of privacy guarantees, beyond what encryption alone can achieve."

The paper was presented this week at the International Colloquium on Automata, Languages, and Programming (ICALP 2014) in Copenhagen. Ohrimenko, who recently received her Ph.D. from Brown University and now works at Microsoft Research, co-authored the work with Roberto Tamassia and Eli Upfal, professors of computer science at Brown, and Michael Goodrich from the University of California–Irvine.

Cloud computing is increasing in popularity as more individuals use services like Google Drive and more companies outsource their data to companies like Amazon Web Services. As the amount of data on the cloud grows, so do concerns about keeping it secure. Most cloud service providers encrypt the data they store. Larger companies generally encrypt their own data before sending it to the cloud to protect it not only from hackers but also to keep cloud providers themselves from snooping around in it.

But while encryption renders data files unreadable, it can't hide patterns of data access. Those patterns can be a serious security issue. For example, a service provider—or someone eavesdropping on that provider—might be able to figure out that after accessing files at certain locations on the cloud server, a company tends to come out with a negative earnings report the following week. Eavesdroppers may have no idea what's in those particular files, but they know that it's correlated to negative earnings.

But that's not the only potential security issue.

"The pattern of accessing data could give away some information about what kind of computation we're performing or what kind of program we're running on the data," said Tamassia, chair of the Department of Computer Science.

Some programs have very particular ways in which they access data. By observing those patterns, someone might be able to deduce, for example, that a company seems to be running a program that processes bankruptcy proceedings.

The Melbourne Shuffle aims to hide those patterns by shuffling the location of data on cloud servers. Ohrimenko named it after a dance that originated in Australia, where she did her undergraduate work.

"The contribution of our paper is specifically a novel data shuffling method that is provably secure and computationally more efficient than previous methods," Ohrimenko said.

It works by pulling small chunks of data down from the cloud and placing them in a user's local memory. Once the data is out of view of the server's prying eyes, it's rearranged—shuffled like a deck of cards—and then sent back to the cloud server. By doing this over and over with new blocks of data, the entirety of the data on the cloud is eventually shuffled.

The result is that data accessed in one spot today, may be in a different spot tomorrow. So even when a user accesses the same data over and over, that access pattern looks to the server or an eavesdropper to be essentially random.

"What we do is we obfuscate the access pattern," Tamassia said. "It becomes unfeasible for the cloud provider to figure out what the user is doing."

The researchers envision deploying their shuffle algorithm through a software application or a hardware device that users keep at their location. It could also be deployed in the form of a tamper-proof chip controlled by the user and installed at the data center of the cloud provider.

However it's deployed, the approach has the promise of lowering the cost of strong data security in an increasingly cloudy computer world.

Provided by Brown University

APA citation: 'Melbourne Shuffle' secures data in the cloud (2014, July 10) retrieved 20 September 2021 from https://phys.org/news/2014-07-melbourne-shuffle-cloud.html

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.