


Protecting sensitive metadata so it can't be used for surveillance

February 26 2020, by Rob Matheson



In a new metadata-protecting scheme, users send encrypted messages to multiple chains of servers, with each chain mathematically guaranteed to have at least one hacker-free server. Each server decrypts and shuffles the messages in random order, before shooting them to the next server in line. Credit: Massachusetts Institute of Technology

MIT researchers have designed a scalable system that secures the metadata—such as who's corresponding and when—of millions of users in communications networks, to help protect the information against possible state-level surveillance.

Data encryption schemes that protect the content of online communications are prevalent today. Apps like WhatsApp, for instance, use "end-to-end encryption" (E2EE), a scheme that ensures third-party eavesdroppers can't read messages sent by [end users](#).

But most of those schemes overlook metadata, which contains information about who's talking, when the messages are sent, the size of message, and other information. Many times, that's all a government or other hacker needs to know to track an individual. This can be especially dangerous for, say, a government whistleblower or people living in oppressive regimes talking with journalists.

Systems that fully protect user metadata with cryptographic privacy are complex, and they suffer scalability and speed issues that have so far limited their practicality. Some methods can operate quickly but provide much weaker security. In a paper being presented at the USENIX Symposium on Networked Systems Design and Implementation, the MIT researchers describe "XRD" (for Crossroads), a metadata-protection scheme that can handle cryptographic communications from millions of users in minutes, whereas traditional methods with the same level of security would take hours to send everyone's messages.

"There is a huge lack in protection for metadata, which is sometimes very sensitive. The fact that I'm sending someone a message at all is not protected by encryption," says first author Albert Kwon Ph.D. '19, a recent graduate from the Computer Science and Artificial Intelligence Laboratory (CSAIL). "Encryption can protect content well. But how can we fully protect users from metadata leaks that a state-level adversary

can leverage?"

Joining Kwon on the paper are David Lu, an undergraduate in the Department of Electrical Engineering and Computer Science; and Srinivas Devadas, the Edwin Sibley Webster Professor of Electrical Engineering and Computer Science in CSAIL.

New spin on mix nets

Starting in 2013, disclosures of classified information by Edward Snowden revealed widespread global surveillance by the U.S. government. Although the mass collection of metadata by the National Security Agency was subsequently discontinued, in 2014 former director of the NSA and the Central Intelligence Agency Michael Hayden explained that the government can often rely solely on metadata to find the information it's seeking. As it happens, this is right around the time Kwon started his Ph.D. studies.

"That was like a punch to the cryptography and security communities," Kwon says. "That meant encryption wasn't really doing anything to stop spying in that regard."

Kwon spent most of his Ph.D. program focusing on metadata privacy. With XRD, Kwon says he "put a new spin" on a traditional E2EE metadata-protecting scheme, called "mix nets," which was invented decades ago but suffers from scalability issues.

Mix nets use chains of servers, known as mixes, and public-private key encryption. The first server receives encrypted messages from many users and decrypts a single layer of encryption from each message. Then, it shuffles the messages in random order and transmits them to the next server, which does the same thing, and so on down the chain. The last server decrypts the final encryption layer and sends the message to the

target receiver.

Servers only know the identities of the immediate source (the previous server) and immediate destination (the next server). Basically, the shuffling and limited identity information breaks the link between source and destination users, making it very difficult for eavesdroppers to get that information. As long as one server in the chain is "honest"—meaning it follows protocol—metadata is almost always safe.

However, "active attacks" can occur, in which a malicious server in a mix net tampers with the messages to reveal user sources and destinations. In short, the malicious server can drop messages or modify sending times to create communications patterns that reveal direct links between users.

Some methods add cryptographic proofs between servers to ensure there's been no tampering. These rely on public key cryptography, which is secure, but it's also slow and limits scaling. For XRD, the researchers invented a far more efficient version of the cryptographic proofs, called "aggregate hybrid shuffle," that guarantees servers are receiving and shuffling message correctly, to detect any malicious server activity.

Each server has a secret private key and two shared public keys. Each server must know all the keys to decrypt and shuffle messages. Users encrypt messages in layers, using each server's secret private key in its respective layer. When a server receives messages, it decrypts and shuffles them using one of the public keys combined with its own private key. Then, it uses the second public key to generate a proof confirming that it had, indeed, shuffled every message without dropping or manipulating any. All other servers in the chain use their secret private keys and the other servers' public keys in a way that verifies this proof. If, at any point in the chain, a server doesn't produce the proof or provides an incorrect proof, it's immediately identified as malicious.

This relies on a clever combination of the popular public key scheme with one called "authenticated encryption," which uses only private keys but is very quick at generating and verifying the proofs. In this way, XRD achieves tight security from public key encryption while running quickly and efficiently.

To further boost efficiency, they split the servers into multiple chains and divide their use among users. (This is another traditional technique they improved upon.) Using some statistical techniques, they estimate how many servers in each chain could be malicious, based on IP addresses and other information. From that, they calculate how many servers need to be in each chain to guarantee there's at least one honest server. Then, they divide the users into groups that send duplicate messages to multiple, random chains, which further protects their privacy while speeding things up.

Getting to real-time

In computer simulations of activity from 2 million users sending messages on a network of 100 servers, XRD was able to get everyone's messages through in about four minutes. Traditional systems using the same server and user numbers, and providing the same cryptographic security, took one to two hours.

"This seems slow in terms of absolute speed in today's communication world," Kwon says. "But it's important to keep in mind that the fastest systems right now [for metadata protection] take hours, whereas ours takes minutes."

Next, the researchers hope to make the network more robust to few users and in instances where servers go offline in the midst of operations, and to speed things up. "Four minutes is acceptable for sensitive messages and emails where two parties' lives are in danger, but

it's not as natural as today's internet," Kwon says. "We want to get to the point where we're sending metadata-protected messages in near real-time."

More information: XRD: Scalable Messaging System with Cryptographic Privacy. www.usenix.org/system/files/ns...ring_kwon_prepub.pdf

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Protecting sensitive metadata so it can't be used for surveillance (2020, February 26) retrieved 9 April 2024 from <https://techxplore.com/news/2020-02-sensitive-metadata-surveillance.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.
