

Secure computers aren't so secure

October 30 2009, by Larry Hardesty



(PhysOrg.com) -- Even well-defended computers can leak shocking amounts of private data. MIT researchers seek out exotic attacks in order to shut them down.

You may update your antivirus software religiously, immediately download all new Windows security patches, and refuse to click any e-mail links ostensibly sent by your bank, but even if your computer is running exactly the way it's supposed to, a motivated attacker can still glean a shocking amount of private information from it. The time it takes to store data in memory, fluctuations in [power consumption](#), even the sounds your computer makes can betray its secrets. MIT researchers centered at the Computer Science and [Artificial Intelligence](#) Lab's Cryptography and Information Security Group (CIS) study such subtle

security holes and how to close them.

In 2005, Eran Tromer, now a postdoc at CIS, and colleagues at the Weizmann Institute in Rehovot, Israel, showed that without any breach of security in the ordinary sense, a seemingly harmless computer program could eavesdrop on other programs and steal the type of secret cryptographic key used by one of the most common Internet encryption schemes. Armed with the key, an attacker could steal a computer user's credit card number, bank account password — whatever the encryption scheme was invoked to protect.

Computer operating systems are supposed to prevent any given program from looking at the data stored by another. But when two programs are running at the same time, they sometimes end up sharing the same cache — a small allotment of high-speed memory where the operating system stores frequently used information. Tromer and his colleagues showed that simply by measuring how long it took to store data at a number of different cache locations, a [malicious program](#) could determine how frequently a cryptographic system was using those same locations. “The memory access patterns — that is, which memory addresses are accessed — are heavily influenced by the specific secret key being used in that operation,” Tromer says. “We demonstrated a concise and efficient procedure for learning the secret keys given just this crude information about the memory access patterns.” Complete extraction of the private key, Tromer says, “takes merely seconds, and the measurements that are needed, of the actual cryptographic process being attacked, can be carried out in milliseconds.”

The encryption system that Tromer was attacking, called AES, was particularly vulnerable because it used tables of precalculated values as a computational short cut, so that encoding and decoding messages wouldn't be prohibitively time consuming. Since Tromer and his colleagues published their results, Intel has added hardware support for

AES to its chips, so that Internet encryption software won't have to rely on such "lookup tables."

In a statement, Intel told the MIT News Office that its decision "was mainly motivated by the performance/efficiency benefits achieved," but that "in addition, there is a potential security benefit since these new instructions can mitigate the possibility of software side channel attacks on AES that have been described in research papers, including those discovered by Tromer, Percival, and Bernstein."

"I think it's fair to say that it's a direct response to the cache-timing attacks against AES," Pankaj Rohatgi, director of hardware security at the data security firm Cryptography Research, says of Intel's move.

Together with CIS cofounder Ron Rivest and CSAIL's Saman Amarasinghe, Tromer is trying to develop further techniques for thwarting cache attacks by disrupting the correlations between [encryption](#) keys and memory access patterns. A couple weeks ago, at the Association for Computing Machinery's Symposium on [Operating Systems](#) Principles, the researchers announced that they had a "proof-of-concept prototype" of a defense system, but they plan to continue testing and refining it before publishing any papers.

Tromer has also been investigating whether cloud computing — the subcontracting of computational tasks to networked servers maintained by companies like Amazon and Google — is susceptible to cache attacks. Many web sites rely on cloud computing to handle sudden surges in their popularity: renting added server space for a few hours at a time can be much cheaper than maintaining large banks of proprietary servers that frequently stand idle.

The word "cloud" is supposed to suggest that this vast agglomeration of computing power is amorphous and constantly shifting, but Tromer and

colleagues at the University of California, San Diego, were able to load their eavesdropping software onto precisely the same servers that were hosting websites they'd targeted in advance. In part, their approach involved spreading their software across a number of servers, then assailing a targeted website with traffic. By spying on the caches of the servers hosting their software, they could determine which were also trying to keep pace with their fake traffic spikes. Once they'd identified the target site's servers, they could use cache monitoring to try to steal secrets.

“Imagine a stock broker that specializes in a specific company,” Tromer says. “If you observe that his virtual machine is particularly active, that could be valuable information. Or you may want to know how popular your competitors' website is. We've actually demonstrated that we can very robustly estimate web server popularity.”

“This has sparked the imagination of both the research community and industry,” Rohatgi says. “I interact with a lot of people in industry, and when they say, ‘Give me the technical basis for this,’ I point to [Tromer and colleagues'] papers.”

Finally, Tromer is continuing work he began as a graduate student, on the use of a “hundred-dollar commodity microphone” to record the very sounds emitted by a computer and analyze them for information about cryptographic keys. So far, Tromer hasn't been able to demonstrate complete key extraction, but he believes he's getting close.

Any information at all about a [computer](#)'s internal workings “is actually fairly damaging,” Rohatgi says. “In some sense, some of these cryptographic algorithms are fairly brittle, and with a little extra information, you can break them.”

Provided by Massachusetts Institute of Technology ([news](#) : [web](#))

Citation: Secure computers aren't so secure (2009, October 30) retrieved 19 April 2024 from <https://phys.org/news/2009-10-secure-computers-arent-so.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.