New hardware design makes data encryption more secure by disguising cloud servers' memory-access patterns

July 2 2013, by Larry Hardesty

Credit: CHRISTINE DANILOFF/MIT

Cloud computing—outsourcing computational tasks over the Internet—could give home-computer users unprecedented processing power and let small companies launch sophisticated Web services without building massive server farms.

But it also raises privacy concerns. A bank of cloud servers could be running applications for 1,000 customers at once; unbeknownst to the

hosting service, one of those applications might have no purpose other than spying on the other 999.

Encryption could make cloud servers more secure. Only when the data is actually being processed would it be decrypted; the results of any computations would be re-encrypted before they're sent off-chip.

In the last 10 years or so, however, it's become clear that even when a computer is handling encrypted data, its memory-access patterns—the frequency with which it stores and accesses data at different memory addresses—can betray a shocking amount of private information.

At the International Symposium on Computer Architecture in June, MIT researchers described a new type of secure hardware component, dubbed Ascend, that would disguise a server's memory-access patterns, making it impossible for an attacker to infer anything about the data being stored. Ascend also thwarts another type of attack, known as a timing attack, which attempts to infer information from the amount of time that computations take.

Computational trade-off

Similar designs have been proposed in the past, but they've generally traded too much computational overhead for security. "This is the first time that any hardware design has been proposed—it hasn't been built yet—that would give you this level of security while only having about a factor of three or four overhead in performance," says Srini Devadas, the Edwin Sibley Webster Professor of Electrical Engineering and Computer Science, whose group developed the new system. "People would have thought it would be a factor of 100."

The "trivial way" of obscuring memory-access patterns, Devadas explains, would be to request data from every address in the

memory—whether a memory chip or a hard drive—and throw out everything except the data stored at the one address of interest. But that would be much too time-consuming to be practical.

What Devadas and his collaborators—graduate students Ling Ren, Xiangyao Yu and Christopher Fletcher, and research scientist Marten van Dijk—do instead is to arrange memory addresses in a data structure known as a "tree." A family tree is a familiar example of a tree, in which each "node" (in this example, a person's name) is attached to only one node above it (the node representing the person's parents) but may connect to several nodes below it (the person's children).

With Ascend, addresses are assigned to nodes randomly. Every node lies along some "path," or route through the tree, that starts at the top and passes from node to node, without backtracking, until arriving at a node with no further connections. When the processor requires data from a particular address, it sends requests to all the addresses in a path that includes the one it's really after.

To prevent an attacker from inferring anything from sequences of memory access, every time Ascend accesses a particular memory address, it randomly swaps that address with one stored somewhere else in the tree. As a consequence, accessing a single address multiple times will very rarely require traversing the same path.

Less computation to disguise an address

By confining its dummy requests to a single path, rather than sending them to every address in memory, Ascend exponentially reduces the amount of computation required to disguise an address. In a separate paper, which is as-yet unpublished but has been posted online, the researchers prove that querying paths provides just as much security as querying every address in memory would.

Ascend also protects against timing attacks. Suppose that the computation being outsourced to the cloud is the mammoth task of comparing a surveillance photo of a criminal suspect to random photos on the Web. The surveillance photo itself would be encrypted, and thus secure from prying eyes. But spyware in the cloud could still deduce what public photos it was being compared to. And the time the comparisons take could indicate something about the source photos: Photos of obviously different people could be easy to rule out, but photos of very similar people might take longer to distinguish.

So Ascend's memory-access scheme has one final wrinkle: It sends requests to memory at regular intervals—even when the processor is busy and requires no new data. That way, attackers can't tell how long any given computation is taking.

The paper is titled "Design space exploration and optimization of path oblivious RAM in secure processors."

More information: Paper (PDF): "Design space exploration and optimization of path oblivious RAM in secure processors"

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: New hardware design makes data encryption more secure by disguising cloud servers' memory-access patterns (2013, July 2) retrieved 26 April 2024 from https://phys.org/news/2013-07-hardware-encryption-disguising-cloud-servers.html