BIZ & IT

At death’s door for years, widely used SHA1 function is now dead

Algorithm underpinning Internet security falls to first-known collision attack.

DAN GOODIN – FEB 23, 2017 1:01 PM | 180

Credit: Bob Embleton

TEXT SETTINGS

For more than six years, the SHA1 cryptographic hash function underpinning Internet security has been at death’s door. Now it’s officially dead, thanks to the submission of the first known instance of a fatal exploit known as a “collision.”

Despite more than a decade of warnings about the lack of security of SHA1, the watershed moment comes as the hash function remains widely used. Git, the world’s most widely used system for managing software

development among multiple people, relies on it for data integrity. The GnuPG e-mail encryption program still deems SHA1 safe. And hundreds if not thousands of big-name software packages rely on SHA1 signatures to ensure installation and update files distributed over the Internet haven’t been maliciously altered.

A collision occurs when the two different files or messages produce the same cryptographic hash. The most well-known collision occurred sometime around 2010 against the MD5 hash algorithm, which is even weaker than SHA1. A piece of nation-sponsored espionage malware known as Flame used the attack to hijack the Windows update mechanism Microsoft uses to distribute patches to hundreds of millions of customers. By forging the digital signature used to cryptographically prove the authenticity of Microsoft servers, Flame was able to spread from one infected computer to another inside targeted networks.

Now, researchers have demonstrated a similar type of real-world attack against SHA1, which ironically was widely adopted after the insecurity of MD5 became well-known. The SHA1 collision is documented in a research paper published Thursday. It presents two PDF files that, despite displaying different content, have the same SHA1 hash. The researchers warned that the same technique—which costs as little as $110,000 to carry out on Amazon’s cloud computing platform—could be used to create collisions in GIT file objects or digital certificates.

“Our work shows that it is now practical to find collisions for SHA1 and that thus it is not secure to use for digital signatures, file integrity, and file identification purposes,” Marc Stevens, the lead researcher, told Ars. “Everyone should migrate to safe standards before real-world attacks happen, not after. Note that attacks can only get better and faster, computational power only becomes cheaper, and attackers have the uncanny ability to be more creative in exploiting vulnerabilities than common expectations.”

RIP

Cryptographers refer to the attack disclosed Thursday as an “identical-prefix” collision, meaning it allows the attacker to create two distinct messages that have the same hash value. This variety is less powerful than the “chosen-prefix” MD5 collision carried out by Flame. In the latter case, attackers can target one or more existing files, such as the digital certificate that a company uses to authenticate its update mechanism. Despite the collision against SHA1 being less powerful, cryptography experts said any real-world identical-prefix attack represented a game-over event for a hashing function.

“In crypto we have the idea that hash function collisions should be really hard to find, even if they’re ‘useless,’” said Johns Hopkins University professor Matt Green, speaking generally about collisions before he learned the specifics of the new SHA1 attack. A real-world collision attack “is the equivalent of finding out that your scalpel wasn’t sterilized properly. It may not verifiably have germs on it, but the whole instrument is considered unsafe.”

Stevens, the lead researcher, agreed. In an e-mail he wrote:

> As an example of the insecurity of SHA1 [the identical-prefix collision] is very meaningful. Chosen-prefix collisions are more powerful and relieve an attacker from some constraints, but identical-prefix collisions can be just as dangerous with a very creative attacker. It is hard to say which things you can’t do with an identical-prefix collision.
>
> But we do know quite some demonstrated threats using identical-prefix collisions: certificates with identical names and different pubkeys, PDF files, PostScript files, TIFF files, JPEG files, Word files, file archives, signed software, Email message/attachment PGP/GPG signatures.

As an example, a landlord could use two colliding rental agreements to trick a prospective tenant into digitally signing a low-rent contract. The landlord could later claim the tenant signed a contract agreeing to a much higher rental price. Fortunately, certificates to HTTPS-protected websites aren’t likely to be affected. Since the beginning of this year, browser-trusted certificate authorities have been barred from relying on SHA1 to sign TLS certificates they issue.

The research is the result of a more than two-year collaboration between researchers at the Centrum Wiskunde & Informatica in the Netherlands and Google’s research security, privacy, and anti-abuse group. The first phase of the attack was run on a heterogeneous CPU cluster that was hosted by Google and spread over eight physical locations. A second and more expensive phase was run on a heterogeneous cluster of K20, K40, and K80 GPUs that were also hosted by Google. Had the researchers performed their attack on Amazon’s Web Services platform, it would have cost $560,000 at normal pricing. Had the researchers been patient and waited to run their attack during off-peak hours, the same collision would have cost $110,000. That’s within the $75,000 to $120,000 range CWI’s Stevens projected in late 2015.

Consistent with Google’s security disclosure policy, the source code for performing the collision attack will be published in 90 days. That means Git and an unknown number of other widely used services that rely on SHA1 have three months to wean themselves and their users off the insecure function. The best candidates for replacements are the SHA256 and SHA3 functions. In the meantime, the researchers have released a tool that detects if files are part of a collision attack. The tool and much more information are available at shattered.io.

Listing image: Bob Embleton

DAN GOODIN *SENIOR SECURITY EDITOR*

Dan Goodin is Senior Security Editor at Ars Technica, where he oversees coverage of malware, computer espionage, botnets, hardware hacking, encryption, and passwords. In his spare time, he enjoys gardening, cooking, and following the independent music scene. Dan is based in San Francisco. Follow him at here on Mastodon and here on Bluesky. Contact him on Signal at DanArs.82.

180 COMMENTS

PREV STORY NEXT STORY

MOST READ

1. Zillow loses thousands of listings in fight over “hidden” homes
2. Ground system issue scrubs first launch of SpaceX's Starship V3 rocket
3. The Internet can't stop watching Figure AI's humanoid robots handling packages
4. Uh-oh, the International Space Station is leaking again
5. Doctors outraged after RFK Jr. fires leaders of key preventive medicine panel

MORE FROM ARS

ABOUT US
STAFF DIRECTORY
ARS NEWSLETTERS
GENERAL FAQ
POSTING GUIDELINES
AI POLICY
RSS FEEDS

CONTACT

CONTACT US
ADVERTISE WITH US
REPRINTS

Manage Preferences