

▶ **SHA1 Cryptographic Hash Update**

A Perspective on Practical Security 2005
by Landon Curt Noll

© Copyright 2005 SystemExperts Corporation. All rights reserved.

▶ SHA1 Cryptographic Hash Update

Overview

On February 13th, 2005, a three-page announcement entitled “Collision Search Attacks on SHA1” was published by Dr. Xiaoyun Wang, and Dr. Hongbo Yu of Shandong University in China and Yiqun Lisa Yin, an independent security consultant. Their announcement summarized the results of their new collision search attacks on the SHA1 cryptographic hash. They also updated their successful attacks against the SHA0 and MD5 cryptographic hashes.

On August 19th, 2005, at the Crypto 2005 conference in Santa Barbara California, it was announced that Dr. Wang improved her February 13th results. The impact of her recent work was to reduce the work needed to successfully attack SHA1 by a factor of 64. On the average, it now requires 10 376 293 541 461 622 784 operations to find a SHA1 collision. While this seems like a huge number, distributed searches using many computers across the Internet have solved problems that were twice as complex. In other words, there exist large collections of computers distributed over the Internet that are capable of finding SHA1 collisions.

On October 31, 2005, the leading cryptographers from around the world met at the US National Institute of Science and Technology (NIST) to discuss these recent developments and what needs to be done in response to these recent discoveries. This paper presents the context of these developments, provides a number of technical details about the current problem, and makes a number of specific near term recommendations about cryptographic hash use.

Context

Cryptographic hashes are mathematical operations that are at the heart of many cryptographic security systems. MD5 and SHA1 are widely used cryptographic hashes. MD5 and SHA1 are commonly used to protect secure web traffic (https/SSL) and encrypted Email (PGP/S/MIME), virtual private networks (VPN, encrypted tunnels), secure remote

access (ssh, sftp), file encryption (encrypted disks, encrypted files), digital certificates, cryptographic document authenticity, and authentication to name just a few. A successful attack on MD5 and SHA1 could compromise the integrity of those facilities.

In this document, we use the term “attack” in the context of quality assurance testing. Dr. Wang’s work is not viewed as aggressive or adversarial by the cryptographic community. Her work was unanimously applauded at the NIST conference as being state of the art. Dr. Wang’s group published their results in order to warn the world about serious weaknesses in the MD5/SHA0/SHA1 hash family.

Technical Analysis

By far, the services that are most vulnerable to the recent attacks are digital signatures and related document authenticity signatures. Other uses of cryptographic hashes such as random mapping, HMACs, and pseudo-random number generation are not as vulnerable to these recent attacks. Nevertheless, it is widely conjectured that once cryptographic collision or a hash is discovered, it is only a matter of time before most other uses of the same cryptographic hash become at risk.

By far the weakest hash family member is the MD5 hash followed by the SHA0 hash. Multiple collisions in the MD5 hash have been discovered and demonstrated. A few SHA0 collisions have also been discovered. To date, no SHA1 collision has been demonstrated. However, as Dr. Wang’s has shown, it is only a matter of time (perhaps a year or two) before a SHA1 collision is discovered. Many cryptographers believe that because SHA224, SHA256, SHA384 and SHA512 are all part of the same hash family, these hash functions will begin to show signs of weakness in 5 years, perhaps less.

Attendees of the NIST conference were reminded to avoid the fallacy of dual hashing. Some people incorrectly believe that the way to overcome MD5 and SHA1 flaws is to compute both the MD5 hash and the SHA1 hash. They operate under the false assumption that combined strength of these two

▶ SHA1 Cryptographic Hash Update

hashes is greater than either hash. It has been shown that this dual hash approach is at best only slightly stronger than SHA1 itself.

The conference addressed a common misconception about Dr. Wang's work. There are some people who operate under the false hope that discovery of two strings that have the same cryptographic hash (hash collision) is mostly a theoretical risk. Some people falsely believe that a known hash collision is of little consequence to any real world security system.

At the conference, it was shown that the discovery of even a single hash collision is a problem for digital signatures and cryptographic document authenticity. For example, one presenter demonstrated how a single MD5 hash collision can be used to construct pairs of digitally signed documents that appear to be dramatically different even though they have the same digital signature. They demonstrated this technique on PostScript, PDF, XML, web pages with JavaScript, web pages with Java, TIFF, and MS Word formatted documents. They speculated on how other common formats such as plain text, plain HTML web pages, web pages with flash, web pages with ActiveX and several other common formats could have their meaning substantially altered and still retain the same cryptographic hash.

A number of proposed modifications to existing hash functions as well as some new hashes were presented at the NIST conference. None of these proposals was viewed as a solution to the current hash problem. A number of the presenters at the NIST conference were even skeptical of their own proposals.

The near unanimous consensus of the NIST conference is that the cryptographic community does not understand how to build cryptographic hashes that are practical, efficient, and resilient to attack. The record of cryptographic hashes is poor. Of the twenty-nine published and peer reviewed cryptographic hashes, fourteen have been shown to be fatally flawed, and seven more show signs of potential flaws. Based on this, the remaining eight should be viewed with some level of skepticism because

they are either too new or because they are too similar to flawed hashes.

The consensus at the NIST conference was that we are at least two years away from being able to establish a process for finding replacement cryptographic hashes. For example, it is conjectured by some that several new hashes may be needed. NIST believes that multiple hashes are needed meet the wide variety of demands placed on hash functions. They conjecture that "one hash for all needs" is neither wise nor practical. It will take at least a year or more to come to a consensus on the number and types of replacement hashed. The consensus is also that we are at least two years away from understanding how to evaluate the quality of a cryptographic hash proposal. Cryptographic designers are at least two years away from incorporating resistance to the recently discovered attacks by Dr. Wang and others.

It is widely believed that NIST won't be able to start the process of establishing an Advanced Hash Standard (AHS) until late 2007 or after. Given how long it took to complete the Advanced Encryption Standard (AES), the AHS process may take three years or more to complete. An AHS may be as far away as late 2010 or more. Once adopted, it will take several years before the AHS is widely deployed in both hardware and software. Some knowledgeable cryptographic skeptics view this timeline as optimistic.

Many believe that the AHS process cannot be a one-time event. Many believe that we will need a new AHS every 6 to 12 years.

Even if cryptographic hash replacements are discovered, tested and accepted as the AHS, there are a number of significant barriers to their deployment. Many of the existing network protocols and message formats are not algorithm agile. They presume the use of MD5 and/or SHA1 and do not allow for the introduction of SHA256 let alone some new hash function. Software interoperability and hardware compatibility are significant barriers to new hash function adoption. A loud and clear call has gone out to the network protocol and information

▶ SHA1 Cryptographic Hash Update

exchange standards bodies developing/modifying standards that can accommodate new hash functions as soon as possible. Moreover, if the AHS process is going to repeat every 6 to 12 years, then standards must allow for the introduction of future hash standards.

By now, you may be wondering if there is any hope left in cryptographic security. The good news is that there is hope. For a start, people are no longer complacent about hash functions. Considerable attention and resources are being brought to bear to find practical solutions as soon as possible. Moreover, while SHA1 is showing signs of significant problems in some areas, SHA1 remains strong in others. Moreover, the SHA256 standard is currently resisting known SHA1 attacks. Theoretical attacks against SHA256 may take a few years to turn into practical attacks

Recommendations for 2005-2006

Based on information presented at the NIST hash conference as well as from other publicly available sources, SystemExperts makes the following recommendations for cryptographic hash use through the end of 2006:

1. Do not use non-standard cryptographic hashes, no matter how strong their proponents claim that they are. The only US cryptographic standard hashes are SHA1, SHA224, SHA256, SHA384, and SHA512.
2. The MD5 hash should not be used. Software and hardware that use MD5 should be removed from service and replaced with equivalent systems that use SHA1 and/or SHA256 as soon as possible. See #3 for both an exception to the rapid removal of MD5. See #4 and #5 for guidelines on how to choose SHA1 and/or SHA256.
3. The one possible exception to immediate elimination of MD5 is MD5 hashed passwords (such as /etc/shadow passwords) and MD5 HMAC. MD5 hashed passwords and MD5 HMAC should be phased out before the end of 2007, if not sooner. They should be replaced with SHA1 and/or SHA256 hash passwords and HMAC. See #4 and #5 for guidelines on how to choose SHA1 and/or SHA256.
4. MD5 hashed passwords should never be transmitted in the clear. One should treat the disclosure of a MD5 hashed password as equivalent to having disclosed the password itself. Until MD5 hashed passwords are replaced, they should only be transmitted over appropriately encrypted links. All backups containing MD5 hashed passwords should be appropriately encrypted. Access to files containing MD5 hashed passwords should be restricted to only super-privileged users and processes. As soon as possible, you should change any password that may have been transmitted in the clear, written to a non-encrypted backup, or accessed by a user or process that is not superprivileged.
5. Existing applications that use SHA1, where possible, should be changed to use SHA256 before the end of 2008, if not sooner. For interoperability with older applications and hardware, these applications may have to also support SHA1. If so, then extra special care must be taken to prevent a “man in the middle” attack from downgrading a SHA256 session into a more vulnerable SHA1 session.
6. Until a new Advanced Hash Standard (AHS) is adopted, all new applications and hardware should be designed to use SHA256. For interoperability with older applications and hardware, these new applications may have to also support SHA1. If so, then extra special care must be taken to prevent a “man in the middle” attack from downgrading a SHA256 session into a more vulnerable SHA1 session.
7. All new applications and protocols must be designed to be algorithm agile. These applications and protocols need to be able to accommodate new hash standards as they are developed. Extra special care must be taken to prevent a “man in the middle” attack from downgrading to either a weaker hash or an unused / unknown hash.

▶ SHA1 Cryptographic Hash Update

8. Existing applications and protocols should be modified to be algorithm agile by the end of 2008, if not sooner. These applications and protocols need to be able to accommodate new hash standards as they are developed. Extra special care must be taken to prevent a “man in the middle” attack from downgrading to either a weaker hash or an unused / unknown hash.
9. SHA384 or SHA512 may be used in place of SHA256 in the above examples. Keep in mind that these hashes are slower than SHA256, which in turn is slower than SHA1, which in turn is slower than MD5. Keep also in mind the need for interoperability, backwards compatibility, and the previously mentioned warnings about “man in the middle” attacks.
10. Because it is possible that SHA1 will become unacceptably weak before 2008, and because SHA256 may become vulnerable to attack before an Advanced Hash Standard (AHS) is adopted, a defense-in-depth approach must be taken. Keep in mind these principles:
 - ▶ Mitigate (reduce and contain) those attacks that you cannot immediately prevent
 - ▶ Detect (monitor and log) those attacks that you cannot immediately mitigate
 - ▶ Audit (after the fact) those attacks that you cannot immediately detect

While these recommendations will not eliminate your cryptographic hash attacks risk, they will help you to avoid, mitigate, or respond to them.

How We Can Help

SystemExperts is continually evaluating the research on hash vulnerabilities as well as progress in the development of new cryptographic hashes.

SystemExperts can help evaluate your current application, application framework, and network protocol use. We can advise you on how much risk you may have in your present day use of digital security, and finally, we can make recommendations about whether you need to change your cryptographic hash solution.