Subject: OFFICIAL COMMENT: MD6
From: Douglas Held <dheld@fortify.com>
Date: Mon, 9 Feb 2009 23:40:45 +0000
To: hash-function@nist.gov
CC: hash-forum@nist.gov

Hello,

The MD6 team amended the submission last month, according to http://groups.csail.mit.edu/cis/md6/.

The NIST submissions page http://csrc.nist.gov/groups/ST/hash/sha-3/Round1/submissions_rnd1.html is still providing the old version of the code.

Am I looking in the correct place? If so, please advise when the update will be available on the NIST site.

Regards,
Douglas Held
Fortify Software

Subject: Re: OFFICIAL COMMENT: MD6
From: Larry Bassham <lbassham@nist.gov>
Date: Tue, 10 Feb 2009 11:08:05 -0500
To: Douglas Held <dheld@fortify.com>

(Same note about "OFFICIAL COMMENT" applies to this one.) The updates will be posted within a day or two.

Larry Bassham

On Feb 9, 2009, at 6:40 PM, Douglas Held wrote:

> Hello,
>
> The MD6 team amended the submission last month, according to http://groups.csail.mit.edu/cis/md6/.
>
> The NIST submissions page http://csrc.nist.gov/groups/ST/hash/sha-3/Round1/submissions_rnd1.html is still providing the old version of the code.
>
> Am I looking in the correct place? If so, please advise when the update will be available on the NIST site.
>
> Regards,
> Douglas Held
> Fortify Software

From: hash-forum@nist.gov on behalf of Ronald L. Rivest [rivest@mit.edu]
Sent: Wednesday, July 01, 2009 10:57 AM
To: Multiple recipients of list
Subject: OFFICIAL COMMENT: MD6

This note is in reply to NIST's request for information regarding tweaks and speed-ups for SHA-3 candidates.

We suggest that MD6 is not yet ready for the next SHA-3 round, and we also provide some suggestions for NIST as the contest moves forward.

(1) NIST has stated that to be competitive, a SHA-3 candidate really needs to be at least as fast as the existing SHA-2 algorithms on the standard reference platforms.

(2) NIST has asked for submitters to provide information regarding tweaks and speedups that submitters would like to make to their algorithms, should their algorithm be chosen for the next SHA-3 round. In particular, submitters should indicate how their algorithm could be made at least as fast as the SHA-2 algorithms, if it is not already that fast.

(3) The submitted algorithm MD6 would need significant speed-up in order to match the SHA-2 speeds on the standard reference platforms. The number of internal rounds in the MD6 compression function would need to be reduced from the current range of 80--168 down to 30--40 or so.

(4) The MD6 submitters feel that it is extremely important that the final SHA-3 algorithm be provably resistant to differential attacks. Indeed, it is the surprising power of differential attacks that stimulated the entire SHA-3 competition. The state of the art is capable of providing such proofs, and NIST should insist that SHA-3 candidates for the next round come supplied with such proofs.

(5) The MD6 team has worked hard to see if a reduced-round version of MD6 could be proven resistant to differential attacks. So far, we have failed to do so.

(6) We have also considered various "tweaks" to the MD6 algorithm that not only reduce the number of rounds, but also change some of the operations in each compression function round. So far, these studies have not yielded a tweaked reduced-round version of MD6 that we can prove is resistant to differential attacks.

(7) Our investigations have also turned up a gap in the proof that the submitted version of MD6 is resistant to differential attacks. (There was a bug in the computer-generated portion of this proof.) We are working to repair this gap, but it seems unlikely that such a repair will help at all with a reduced-round version of MD6.

Thus, while MD6 appears to be a robust and secure cryptographic hash algorithm, and has much merit for multi-core processors, our inability to provide a proof of security for a reduced-round (and possibly tweaked) version of MD6 against differential attacks suggests that MD6 is not ready for consideration for the next SHA-3 round.

We are continuing to work on MD6, and this situation may change.

However, we are at this stage not particularly optimistic.

Going forward, our suggestions to NIST for the next round include (in addition to the above-noted requirement for provable resistance to differential attacks):

(a) Make sure that SHA-3 includes a "tree-hashing" mode that is suitable for use on multi-core processors. This shouldn't be just a hand-wave saying that "of course, any sequential hash algorithm can be adapted for use in tree-hashing mode", but should be a detailed spec with the i's dotted and t's crossed.

(b) Make sure that the cost of implementing SHA-3 in "clean" mode is well understood. Here "clean mode" means "in a manner that provides resistance to timing attacks", e.g. by avoiding instruction with data-dependent timings or data-dependent memory usage patterns. Each submitted algorithm should include not only optimized implementations for the standard reference platforms, but also optimized "clean" implementations. The speeds of such "clean" implementations should be considered in the final SHA-3 selection.

We hope this clarification of the current status of MD6, and these suggestions to NIST, will help the SHA-3 process reach the best possible result.

Ronald L. Rivest
for the entire MD6 team (Ben Agre, Dan Bailey, Chris Crutchfield, Yevgeniy Dodis, Kermit Elliot Fleming, Asif Khan, Jayant Krishnamurthy, Yuncheng Lin, Leo Reyzin, Emily Shen, Jim Sukha, Drew Sutherland, Eran Tromer, Yiqun Lisa Yin).

--

Ronald L. Rivest
Room 32-G692, Stata Center, MIT, Cambridge MA 02139
Tel 617-253-5880, Email <rivest@mit.edu>