

Brute-Force Approach to Crack MD5 Hashes with MPI

Background Information

- Passwords remain the single most common point of failure in system security
- MD5 was created by Ron Rivest in 1991 and was widely used in the 90's to encrypt SSL sessions and passwords
- Analytic weaknesses in the algorithm was discovered in 1996 and 2004
- It has been deprecated in favor of the SHA family of algorithms

The Algorithm



The algorithm input is a variable length string, the output is a 32-char length hexadecimal hash. Code was taken directly from the creators home page at...

<http://people.csail.mit.edu/rivest/md5.c>

Psuedo-Code Concept

original string → MD5 routine → hash
test string → MD5 routine → hash2

The program generates test strings to run through the MD5 algorithm and compares the output to the hash to crack. Once hash2 = hash, the test string is outputted to screen and the program exits.

Parallel Approach

Divide the “a-z”, “A-Z”, and “0-9” at each character length (except 1) by the number of processors and parse out to the nodes, e.g. 2 char test work division = aa-mz @ node 0, na-zz @ node 1 for a two-node system.

Theoretical Times

Mathematically, the number of combinations to crack an N-char string on average is $(N/2)^C$ where C is number characters in the set (26 for lowercase alpha). Assuming 40e6 comps/sec...

3-char	.009s
4-char	1.67s
5-char	9m25s
6-char	17.65 hrs

Observed Serial Times

Dual-Core Ubuntu AMI on the AWS

csi	.065s
abba	1.071s
hello	8m38.646s
nvidia	Quit at the 3 hour mark.

MPI Time Trials

6-character initial string (“nvidia”)

cores real time

8 35.541s

6 35.828s

4 39.017s

2 48.586s

5-character initial string (“hello”)

cores real time

8 2.282s

6 2.482s

4 4.925s

2 10.943s

MPI Time Trials

4-character initial string (“abba”)

8 .336s

6 .295s

4 .273s

2 .232s

3-character initial string (“csi”)


8 .337s

6 .293s

4 .274s

2 .207s

AWS Time Trials



Development Problems

With the limited availability of the GMICE cluster, the code was run on the Amazon Cloud with an Ubuntu 64-Bit 8-Core AMI instance. The amount of real time allowed using all 8 cores (-np 8) was limited to 1m30s while running the code. This limitation cut down the type and amount of characters that could be searched to 6 characters lower-case only in order to do timings.