

(cuda)Hashcat

Hashcat

„Installation“ unter Apple Macintosh OS X 10.11.x:

```
$ cd Downloads
$ wget http://hashcat.net/files/hashcat-2.00.7z
[...]
$ 7z x hashcat-2.00.7z
[...]
$ cd hashcat-2.00
```

SHA512 (Linux)

Welche Benutzer gibt es auf dem System (Debian GNU/Linux)?

```
debian$ cat /etc/passwd
[...]
everybody:x:1001:1001:Hashcat Test Account,,,:/home/everybody:/bin/bash
[...]
```

Wie werden die Passwörter auf dem System „gehasht“?

```
debian$ grep -E "^ENCRYPT_METHOD" /etc/login.defs
ENCRYPT_METHOD SHA512
```

Den „anzugreifenden“ Hash kopieren wir in eine Datei:

```
debian# grep -F "everybody" /etc/shadow | cut -d ":" -f 2 | tee
~/Downloads/everybody_shadow.sha512
$6$qzwwrTUI$a079fjxzggxBezWq8fvUrKH20XiR5Y/VTKoMsJ9WXjbo7WZWMLbDYlamkwjoIV/N
G5WdoYN0RIPtIdNW6yLZa.
```

Welches Hashcat-„Modul“ kann solche Hashs „angreifen“?

```
osx$ ./hashcat-cli64.app --help | grep -F -i 'sha' | grep -F -i '512' | grep
-F -i 'unix'
1800 = SHA-512(Unix)
```

Ich weiß „nur“, dass es sich um ein vierstelliges Passwort aus Kleinbuchstaben handelt (die Datei mit dem Hash habe ich bereits in das Verzeichnis von Hashcat kopiert!):

```
osx$ ./hashcat-cli64.app -m 1800 -a 3 everybody_shadow.sha512 ?l?l?l?l
Initializing hashcat v2.00 with 8 threads and 32mb segment-size...
```

```
Added hashes from file shadow_sha512.txt: 1 (1 salts)
```

```
Activating quick-digest mode for single-hash with salt
```

```
[s]tatus [p]ause [r]esume [b]ypass [q]uit => s
```

```
Input.Mode: Mask (?l?l?l?l) [4]
Index.....: 0/1 (segment), 456976 (words), 0 (bytes)
Recovered.: 0/1 hashes, 0/1 salts
Speed/sec.: - plains, 1.10k words
Progress..: 14200/456976 (3.11%)
Running...: 00:00:00:13
Estimated.: 00:00:06:43
```

```
$6$qzwwrTUI$a079fjxzggxBezWq8fvUrKH20XiR5Y/VTKoMsJ9WXjbo7WZWMLbDYlamkwjoIV/N
G5WdoYN0RIPtIdNW6yLZa.:nstl
All hashes have been recovered
```

```
Input.Mode: Mask (?l?l?l?l) [4]
Index.....: 0/1 (segment), 456976 (words), 0 (bytes)
Recovered.: 1/1 hashes, 1/1 salts
Speed/sec.: - plains, 1.08k words
Progress..: 282948/456976 (61.92%)
Running...: 00:00:04:22
Estimated.: 00:00:02:41
```

```
Started: Sat Jan 16 11:04:28 2016
Stopped: Sat Jan 16 11:08:50 2016
```

Wenn ich alle Zeichen (nicht „nur“ Kleinbuchstaben) berücksichtigen möchte, dauert das ganze schon erheblich länger:

```
osx$ ./hashcat-cli64.app -a 3 -m 1800 everybody_shadow.sha512 ?a?a?a?a
Initializing hashcat v2.00 with 8 threads and 32mb segment-size...
```

```
Added hashes from file shadow_sha512.txt: 1 (1 salts)
Activating quick-digest mode for single-hash with salt
```

```
[s]tatus [p]ause [r]esume [b]ypass [q]uit => s
```

```
Input.Mode: Mask (?a?a?a?a) [4]
Index.....: 0/1 (segment), 81450625 (words), 0 (bytes)
Recovered.: 0/1 hashes, 0/1 salts
Speed/sec.: - plains, 1.09k words
Progress..: 13912/81450625 (0.02%)
Running...: 00:00:00:13
Estimated.: 00:20:44:04
```

```
[s]tatus [p]ause [r]esume [b]ypass [q]uit => q
```

```
[...]
```

NTLM (Windows)

Wenn „die“ Festplatte mit „der“ Windows-Installation unter Debian GNU/Linux 8.x unter „/media/user/222444A1244479B5“ automatisch eingehängt wurde, kopieren wir den NTLM-Hash eines bestimmten Benutzers wie folgt in eine Datei:

```
debian$ sudo aptitude install samdump2
[...]
```

```
debian$ cd /media/user/222444A1244479B5/Windows/System32/config
debian$ samdump2 SYSTEM SAM | grep -F "everybody" | cut -d ":" -f 4 | tee
~/Downloads/hashcat-2.00/everybody_SAM.ntlm
1ea1fd7b4931ec9255cda7cb6060b092
```

In dem selben 7z-Archiv (siehe oben) sind auch Binaries für Linux enthalten (die Installation ist vergleichbar einfach!):

```
debian$ ./hashcat-cli64.bin -a 3 -m 1000 everybody_SAM.ntlm ?a?a?a?a
Initializing hashcat v2.00 with 8 threads and 32mb segment-size...
```

```
Added hashes from file ../samdump2_modified.txt: 1 (1 salts)
Activating quick-digest mode for single-hash
```

```
[s]tatus [p]ause [r]esume [b]ypass [q]uit =>
```

```
Input.Mode: Mask (?a?a?a?a) [4]
Index.....: 0/1 (segment), 81450625 (words), 0 (bytes)
Recovered.: 0/1 hashes, 0/1 salts
Speed/sec.: 95.62M plains, 95.62M words
Progress..: 81450625/81450625 (100.00%)
Running...: --:--:--:--
Estimated.: --:--:--:--
```

```
Started: Tue Jan 19 15:33:35 2016
```

```
Stopped: Tue Jan 19 15:33:37 2016
```

Hier ist schon mit dem bloßen Auge erkennbar, dass NTLM-Hashes „etwas“ schneller berechnet werden können als SHA512-Hashes!

cudaHashcat

Installation unter Debian GNU/Linux 8.x (mit einem relativ alten Nvidia-Treiber):

```
$ sudo aptitude install libcuda1
[...]
```

```
$ mkdir ~/Downloads
$ cd Downloads
$ wget http://hashcat.net/files/cudaHashcat-2.01.7z
[...]
$ 7z x cudaHashcat-2.01.7z
[...]
$ cd cudaHashcat-2.01
```

SHA512 (Linux)

Analog zum „CPU-Durchgang“ (siehe oben) jetzt mit der (nicht wirklich potenten) GPU inkl. schreiben von Ergebnissen in eine Datei:

```
debian$ ./cudaHashcat64.bin -w 3 -m 1800 -a 3 -o outfile.txt
~/Downloads/shadow_sha512.txt ?l?l?l?l
cudaHashcat v2.01 starting...

Device #1: GeForce 605, 1023MB, 1046Mhz, 1MCU
Device #1: WARNING! Kernel exec timeout is not disabled, it might cause you
errors of code 702

Hashes: 1 hashes; 1 unique digests, 1 unique salts
Bitmaps: 16 bits, 65536 entries, 0x0000ffff mask, 262144 bytes, 5/13 rotates
Applicable Optimizers:
* Zero-Byte
* Single-Hash
* Single-Salt
* Brute-Force
Watchdog: Temperature abort trigger set to 90c
Watchdog: Temperature retain trigger set to 80c
Device #1: Kernel ./kernels/4318/m01800.sm_21.64.cubin
Device #1: Kernel ./kernels/4318/markov_le_v1.sm_21.64.cubin
Device #1: Kernel ./kernels/4318/amp_a3_v1.sm_21.64.cubin

Session.Name...: cudaHashcat
Status.....: Cracked
Input.Mode.....: Mask (?l?l?l?l) [4]
Hash.Target....: $6$qzwwrTUI$ao79fjxzggxBezWq8fvUrKH20XiR5...
Hash.Type.....: sha512crypt, SHA512(Unix)
Time.Started...: Mon Jan 18 12:10:02 2016 (3 mins, 37 secs)
Speed.GPU.#1...:      897 H/s
Recovered.....: 1/1 (100.00%) Digests, 1/1 (100.00%) Salts
Progress.....: 194560/456976 (42.58%)
Rejected.....: 0/194560 (0.00%)
Restore.Point...: 6144/17576 (34.96%)
HWMon.GPU.#1...: -1% Util, 63c Temp, 50% Fan

Started: Mon Jan 18 12:10:02 2016
Stopped: Mon Jan 18 12:13:41 2016
```

Was steht nun in der Datei?

```
$ cat outfile.txt
$6$qzwwrTUI$a079fjxzggxBezWq8fvUrKH20XiR5Y/VTKoMsJ9WXjbo7WZWMLbDYlamkwjoIV/N
G5WdoYN0RIPtIdNW6yLZa.:nstl
```

NTLM (Windows)

Und das ganze mit einem NTLM-Hash:

```
debian$ cd ~/Downloads/cudaHashcat-2.01
debian$ ./cudaHashcat64.bin -w 3 -m 1000 -a 3 -o everybody_SAM.ntlm
everybody_PLAIN.txt ?l?l?l?l
cudaHashcat v2.01 starting...

Device #1: GeForce 605, 1023MB, 1046Mhz, 1MCU
Device #1: WARNING! Kernel exec timeout is not disabled, it might cause you
errors of code 702

Hashes: 1 hashes; 1 unique digests, 1 unique salts
Bitmaps: 16 bits, 65536 entries, 0x0000ffff mask, 262144 bytes, 5/13 rotates
Applicable Optimizers:
* Zero-Byte
* Precompute-Init
* Precompute-Merkle-Demgard
* Meet-In-The-Middle
* Early-Skip
* Not-Salted
* Not-Iterated
* Single-Hash
* Single-Salt
* Brute-Force
* Scalar-Mode
* Raw-Hash
Watchdog: Temperature abort trigger set to 90c
Watchdog: Temperature retain trigger set to 80c
Device #1: Kernel ./kernels/4318/m01000_a3.sm_21.64.cubin
Device #1: Kernel ./kernels/4318/markov_le_v1.sm_21.64.cubin

ATTENTION!
  The wordlist or mask you are using is too small.
  Therefore, oclHashcat is unable to utilize the full parallelization power
of your GPU(s).
  The cracking speed will drop.
  Workaround:
https://hashcat.net/wiki/doku.php?id=frequently\_asked\_questions#how\_to\_create\_more\_work\_for\_full\_speed

INFO: approaching final keyspace, workload adjusted
```

```
Session.Name...: cudaHashcat
Status.....: Cracked
Input.Mode.....: Mask (?l?l?l?l) [4]
Hash.Target....: 1ea1fd7b4931ec9255cda7cb6060b092
Hash.Type.....: NTLM
Time.Started...: 0 secs
Speed.GPU.#1...: 115.0 MH/s
Recovered.....: 1/1 (100.00%) Digests, 1/1 (100.00%) Salts
Progress.....: 456976/456976 (100.00%)
Rejected.....: 0/456976 (0.00%)
HWMon.GPU.#1...: -1% Util, 53c Temp, 40% Fan

Started: Mon Jan 18 15:22:06 2016
Stopped: Mon Jan 18 15:22:07 2016
```

Hier meckert die GPU-Version, da nicht genug Berechnungen parallelisiert werden können!

Lösung: kleine „Wortlisten“ der CPU übergeben und in einer anderen Konsole eine ergänzende größere „Wortliste“ der GPU übergeben!

Haben wir ein Ergebnis?

```
$ cat everybody_PLAIN.txt
1ea1fd7b4931ec9255cda7cb6060b092:nstl
```

Benchmark

MacBook Pro (Retina, 15-inch, Late 2013):

```
$ ./hashcat-cli64.app -b
Initializing hashcat v2.00 with 8 threads and 32mb segment-size...

Device.....: Intel(R) Core(TM) i7-4850HQ CPU @ 2.30GHz
Instruction set..: x86_64
Number of threads: 8

Hash type: MD5
Speed/sec: 85.83M words

Hash type: SHA1
Speed/sec: 46.12M words

Hash type: SHA256
Speed/sec: 25.45M words

Hash type: SHA512
Speed/sec: 5.71M words

Hash type: bcrypt, Blowfish(OpenBSD)
```

Speed/sec: 5.55k words

Hash type: NTLM

Speed/sec: 81.86M words

Hash type: WPA/WPA2

Speed/sec: 3.42k words

FUJITSU ESPRIMO P910:

```
$ ./hashcat-cli64.bin -b
```

```
Initializing hashcat v2.00 with 8 threads and 32mb segment-size...
```

```
Device.....: Intel(R) Core(TM) i7-3770 CPU @ 3.40GHz
```

```
Instruction set.: x86_64
```

```
Number of threads: 8
```

```
Hash type: MD5
```

```
Speed/sec: 88.75M words
```

```
Hash type: SHA1
```

```
Speed/sec: 52.09M words
```

```
Hash type: SHA256
```

```
Speed/sec: 24.96M words
```

```
Hash type: SHA512
```

```
Speed/sec: 8.50M words
```

```
Hash type: bcrypt, Blowfish(OpenBSD)
```

```
Speed/sec: 5.87k words
```

```
Hash type: NTLM
```

```
Speed/sec: 90.71M words
```

```
Hash type: WPA/WPA2
```

```
Speed/sec: 4.42k words
```

FUJITSU ESPRIMO P910 (Teil 2):

```
$ ./cudaHashcat64.bin -b
```

```
cudaHashcat v2.01 starting in benchmark-mode...
```

```
Device #1: GeForce 605, 1023MB, 1046Mhz, 1MCU
```

```
Hashtype: MD5
```

```
Workload: 1024 loops, 256 accel
```

```
Speed.GPU.#1.: 168.7 MH/s
```

```
Hashtype: SHA1
```

Workload: 1024 loops, 256 accel

Speed.GPU.#1.: 41459.3 kH/s

Hashtype: SHA256

Workload: 1024 loops, 256 accel

Speed.GPU.#1.: 19300.3 kH/s

Hashtype: SHA512

Workload: 256 loops, 256 accel

Speed.GPU.#1.: 5210.5 kH/s

Hashtype: bcrypt, Blowfish(OpenBSD)

Workload: 32 loops, 2 accel

Speed.GPU.#1.: 46 H/s

Hashtype: NTLM

Workload: 1024 loops, 256 accel

Speed.GPU.#1.: 251.2 MH/s

Hashtype: WPA/WPA2

Workload: 1024 loops, 32 accel

Speed.GPU.#1.: 2785 H/s

From:

<http://wiki.neumannsland.de/> - **Patricks DokuWiki**

Permanent link:

<http://wiki.neumannsland.de/wip:hashcat>

Last update: **2019/09/20 09:16**

